

# ISCR

## Annual Report

FY1998



Institute for Scientific Computing Research



The University Relations Program (URP) encourages collaborative research between Lawrence Livermore National Laboratory (LLNL) and the University of California campuses. The Institute for Scientific Computing Research (ISCR) actively participates in such collaborative research, and this report details the Fiscal Year 1998 projects jointly served by URP and ISCR. For a full discussion of all URP projects in FY 1998, please request a copy of the URP FY 1998 Annual Report by contacting

Lawrence Livermore National Laboratory  
Christina M. Budwine, University Relations Program  
P. O. Box 808, L-413  
Livermore, CA 94551

**UCRL-LR-133866-98**



#### DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from the Office of Scientific and Technical Information  
P.O. Box 62, Oak Ridge, TN 37831  
Prices available from (423) 576-8401


Available to the public from the National Technical Information Service  
U.S. Department of Commerce  
5285 Port Royal Rd., Springfield, VA 22161



# ISCR Fiscal Year 1998 Report

Advances in scientific computing research have never been more vital to the core missions of Lawrence Livermore National Laboratory than they are today. These advances are evolving so rapidly, and over such a broad front of computational science, that to remain on the leading edge, the Laboratory must collaborate with many academic centers of excellence. In FY 1998, ISCR dramatically expanded its interactions with academia through collaborations, visiting faculty, guests and a seminar series. The pages of this annual report summarize the activities of the 63 faculty members and 34 students who participated in ISCR collaborative activities during FY 1998.

The 1998 ISCR call for proposals issued by the University Collaborative Research Program (UCRP) resulted in eight awards made by the University of California Office of the President to research teams at UC San Diego, UC Davis, UC Los Angeles, and UC Berkeley. These projects are noted below.



John M. Fitzgerald  
Acting Director, ISCR

## **Software Infrastructure for Multi-Tier Implementation of Structured Adaptive Mesh Hierarchies**

Professor Scott Baden, UCSD  
LLNL collaborator, Jonathan May (CASC)

## **A Computational Investigation of Finite-Strain Plasticity Models for Crystalline Solids**

Professor Panayiotis Papadopoulos, UCB  
LLNL collaborator, Daniel J. Nikkel Jr. (ME)

## **Massively Parallel Computation Applied to Finite Element Analysis**

Professor Tony Keaveny, UCB  
LLNL collaborator, Karin Hollerbach (CASC)

## **Computational Fluid Dynamic Studies of Arterial Flow Disturbances Induced by Intravascular Stents**

Professor Abdul Barakat, UCD  
LLNL collaborator, Patrick Fitch (BBRP)

## **Parallel Multigrid Methods on Unstructured Grids for Scientific Computing**

Professor Tony Chan, UCLA  
LLNL collaborator, Robert Falgout (CASC)

## **Parallel Particle-in-Code Modeling of Semiclassical Quantum Models**

Professor Viktor Decyk, UCLA  
Professor John M. Dawson, UCLA  
LLNL collaborator, Dennis Hewett (CASC)

## **Development of a Three-Dimensional Relativistic Particle-In-Cell Code for Studying the Production of Useful Electron Bunches Using Ultra-Intense Laser Pulses**

Prof Warren B. Mori, UCLA  
LLNL collaborator, Scott Wilks (X-Division)

## **Methods of Simulation for Localized Multiply-Scaled Condensed Materials**

Professor John Weare, UCSD  
LLNL collaborator, Scott Kohn (CASC)

ISCR is now part of the Laboratory's Center for Applied Scientific Computing (CASC). Many CASC scientists participate actively in ISCR–University collaborations, as noted. The eight collaborations shown represent innovative research efforts supported by ISCR in FY 1998. Abstracts discussing each of these collaborations begin on page 79.

The Accelerated Strategic Computing Initiative (ASCI) established Academic Strategic Alliances Program (ASAP) centers located at: Stanford University; California Institute of Technology; University of Chicago; University of Utah, Salt Lake; and University of Illinois, Urbana-Champaign.

The ASCI Alliances strategy was established to enhance overall ASCI goals by establishing technical interactions between the Department of Energy, Defense Programs laboratories (Lawrence Livermore, Los Alamos, and Sandia National Laboratories), and leading-edge research-and-development universities in the United States.

ISCR has partnered with the LLNL ASCI Program Office to facilitate these collaborations. In FY 1998, ISCR hosted ASCI Alliances student guests who visited LLNL for collaborations during the summer months.


We hope that you enjoy this report on ISCR's many activities throughout FY 1998. For further information about the Institute, please contact me at your convenience.

**John Fitzgerald**  
Acting Director

Institute for Scientific Computing Research  
Lawrence Livermore National Laboratory  
P.O. Box 808, L-561  
Livermore, CA 94551

fitzgerald1@llnl.gov  
925/ 422-4024

ISCR Fiscal Year 1998 in Review.....	4
FY 1998 Seminar Series .....	11
Sabbatical Visitor .....	67
Subcontracts, ASCI Level 3 .....	71
University Collaborative Research Program Proposals .....	77
LDRD Projects .....	89
Students and Faculty .....	95



## FY 1998 Seminar Series (in reverse chronological order)

Michael Pernice, University of Utah	September 29, 1998
Dean Dager, University of California, Los Angeles	September 28, 1998
Malvin H. Kalos, Cornell University	September 23, 1998
Steven Lumetta, University of Illinois, Urbana-Champaign	September 17, 1998
Leonid Oliker, NASA Ames Research Center	September 11, 1998
Roger L. Haskin, IBM Almaden Research Center	September 3, 1998
Jonathan Dursi, University of Chicago	August 31, 1998
Yuan-Nan Young, University of Chicago	August 31, 1998
Abraham Szoke, Lawrence Livermore National Laboratory	August 25, 1998
Timothy Barth, Numerical Aerospace Simulation, NASA	August 24, 1998
Dave Beazley, University of Chicago	August 21, 1998
Steve Parker, University of Utah	August 19, 1998
Louisa Raschid, University of Maryland	August 15, 1998
Michael Franklin, University of Maryland	August 15, 1998
Harland Glaz, University of Maryland	August 12, 1998
Karen L. Karavanic, University of Wisconsin, Madison	August 7, 1998
Howard Elman, University of Maryland	July 31, 1998
George Bebis, University of Nevada, Reno	July 22, 1998
Chris Brislawn, Los Alamos National Laboratory	July 17, 1998
Andrew Grimshaw, University of Virginia	July 10, 1998
Bernd Hamann, University of California, Davis	July 1, 1998
Raycho Lazarov, Texas A & M University	June 29, 1998
Ulrich Ruede, Universitaet Augsburg	June 25, 1998
Dirk Grunwald, University of Colorado, Boulder	June 19, 1998
Roland Freund, Bell Labs	June 18, 1998
Ian Foster, Argonne National Laboratory	June 16, 1998
Laura Myers Haas, IBM Almaden	June 12, 1998
Dimitri Mavriplis, ICASE, NASA Langley	June 11, 1998
David Keyes, Old Dominion University	May 29, 1998
Hongyan Zha, Pennsylvania State University	May 22, 1998
Jack Dongarra, University of Tennessee	May 19, 1998
Chandrajit L. Bajaj, University of Texas at Austin	May 1, 1998
Dean H. Judson, University of Nevada	April 28, 1998
Paul Dubois, Lawrence Livermore National Laboratory	April 20, 1998
Raghu Machiraju, Mississippi State University	April 17, 1998

Nabil Adam, Rutgers University	April 17, 1998
and Igg Adiwijaya, Rutgers University	
Peter Schroeder, California Institute of Technology	April 14, 1998
Kenneth I. Joy, University of California, Davis	April 10, 1998
Malvin H. Kalos, Cornell University	March 30, 1998
Ann Hollister, Louisiana State Univ. Medical Center	March 27, 1998
Louise Focht, Avanta Orthopedics, Inc.	March 27, 1998
Jinchao Xu, Penn State University	March 20, 1998
Jesse Barlow, Penn State University	March 19, 1998
John Feo, Tera Computers	March 9, 1998
Tamara G. Kolda, Oak Ridge National Laboratory	February 26, 1998
John Bell, Lawrence Berkeley National Laboratory	February 24, 1998
Sam Uvelton, Numerical Aerospace Simulation, NASA	February 20, 1998
Paul Woodward, University of Minnesota	February 20, 1998
Kathy Holian, Los Alamos National Laboratory	February 11, 1998
Dan Quinlan, Los Alamos National Laboratory	February 6, 1998
Michael Holst, University of California, Irvine	January 16, 1998
David Keyes, Old Dominion University	December 19, 1997
David H. Bailey, NASA Ames Research Center	December 16, 1997
James Browne, University of Texas at Austin	December 12, 1997
Thomas F. Russell, University of Colorado, Denver	December 9, 1997
Randy Bank, University of California San Diego	November 11, 1997
Sing-Lok (Justin) Wan, Univ. of California, Los Angeles	November 10, 1997

## Visitors, Guests, Consultants, and Researchers

### Sabbatical Visitor

Raytcho Lazarov, Texas A&M University

### Participating Guests

Marsha Berger, New York University

George Byrne, Illinois Institute of Technology

David Dean, University of Colorado

Harry Dwyer, University of California, Davis

Alejandro Garcia, San Jose State University

Bernd Hamann, University of California, Davis

Stanley Johnson, Lehigh University

Joel Kaiser, University of California, Davis

Martin Lades, Genetrace

David Larson, Bay Area Research Corporation

Gerry Elbridge Puckett, University of California, Davis

Yousef Saad, University of Minnesota

Paul Saylor, University of Illinois, Urbana-Champaign

Jeffrey Scoggs, North Carolina State University

Kenneth Trauner, University of California, Davis Medical School

Daniel Wake, TMA, Inc.

### Consultants

Berni Alder, University of California Professor Emeritus

Randolph Bank, University of California San Diego

Anne Greenbaum, University of Washington

Thomas Manteuffel, University of Colorado

Stephen McCormick, University of Colorado

Linda Petzold, University of California, Santa Barbara

Steve Schaffer, New Mexico Tech

Homer Walker, Worcester Polytechnic Institute

## Visitors

Marian Brezina, University of Colorado  
David Brown, Los Alamos National Laboratory  
Larry Carter, University of California, San Diego  
Angela Cheer, University of California, Davis  
Al Davis, University of Utah  
Harry Dwyer, University of California, Davis  
James Demmel, University of California, Berkeley  
Oliver Duschka, Socratix System  
Raymond Scott Fellers II, University of California, Berkeley  
Ernesto Gomez, University of Chicago  
Charles Hansen, University of Utah  
Thomas Henderson, University of Utah  
Michael Allen Heroux, Sandia National Laboratories Albuquerque  
Chris Johnson, University of Utah  
Joel Kaiser, University of California, Davis  
Thomas Ivan Milac, University of Washington  
Lloyd Minick, Socratix System  
Mark Alan Musan, Stanford University and Socratix System  
Esmond Gee-Ying Ng, Oak Ridge National Laboratory  
Julia Elizabeth Rice, IBM Almaden  
Jennifer M. Schopf, University of California, San Diego  
Andrew P. Strelzoff, University of California, Santa Barbara  
Wei-Pai Tang, University of Waterloo  
Kathleen Thompson, Socratix System  
Angus Taylor, University of California, Berkeley  
Michael Tocci, Worcester Polytechnic Institute

## Visitors, Guests, Consultants, and Researchers (cont.)

### Subcontracts, ASCI Level 3

Bruce Boghosian, Boston University  
David Keyes, Old Dominion University  
John Ruge, Front Range Scientific Computation, Inc.

### University Collaborative Research Program Proposals

Scott Baden, University of California, San Diego  
Abdul Barakat, University of California, Davis  
Tony Chan, University of California, Los Angeles  
Viktor Decyk and John Dawson, University of California, Los Angeles  
Tony Keaveny, University of California, Berkeley  
Warren Mori, University of California, Los Angeles  
Panayiotis Papadopoulos, University of California, Berkeley  
John Weare, University of California, San Diego

### LDRD Projects

Karin Hollerbach, LLNL Center for Applied Scientific Computing  
Peter Brown, LLNL Center for Applied Scientific Computing

### Postdoctoral Researchers

Raymond S. Fellers II  
Madhaven Ganesh  
Brian J. Miller  
Thomas F. Rutaganira



# Students and Faculty

## Department of Applied Science Faculty

Nelson L. Max  
Garry H. Rodrique

## Department of Applied Science Students

Paul E. Covello  
David S. Miller  
Johnathan E. Rochez  
Jay Feltus Thomas

## Visiting Students and Faculty

Mark Adams, University of California, Berkeley  
Igg Adiwijaya, Rutgers University  
Travis Austin, University of Colorado, Boulder  
Lora Ballinger, University of Utah  
George Bebis, University of Nevada, Reno  
Martin Bertram, University of California, Davis  
Kathleen Bonnell, California State University, Dominguez Hills  
Dan Bullok, University of Illinois  
Tim Chartier, University of Colorado, Boulder  
Jonathan Dursi, University of Chicago  
Chris Higginson, University of Colorado, Boulder  
Rithea Hong, Texas A&M University  
Rachel Karchin, Stanford University  
Falko Kuester, University of California, Davis  
Jonathan Pearlman, University of California, Berkeley  
Lyn Reid, University of Washington  
Natasha Sager, Las Positas Junior College  
Michael Wittman, University of California, Davis  
Yuan-Nan Young, University of Chicago

## Students and Faculty (cont.)


### National Physical Science Consortium (NPSC)

Lora Ballinger, University of Utah

Rachel Karchin, Stanford University

Imelda Kirby, University of Arizona

Megan Thomas, University of California, Davis



## FY 1998 Seminar Series

---

(in reverse chronological order)





September 29, 1998

# Hybrid Approaches for Solution of Large-Scale Systems of Nonlinear Equations

*Michael Pernice**University of Utah  
pernice@sneffels.chpc.utah.edu*

## Abstract

There has been considerable discussion of the relative merits of multigrid and Newton-based methods for solving large-scale systems of nonlinear equations. The full approximation scheme (FAS) promises cheap iterations and high rates of convergence, but this is difficult to achieve without the correct combination of interlevel transfers and smoothing strategies. While superlinear rates of convergence are attractive, Newton-based iterations are more expensive (both in terms of operations and storage) and even with globalization strategies performance can be sensitive to the choice of an initial approximation. Recent advances in the development of inexact Newton methods have made this approach more competitive, but a good preconditioner is still necessary to achieve satisfactory performance.

This debate is largely misdirected. Since both approaches require selection of components whose effectiveness is highly problem-dependent, it is unlikely that one strategy will emerge as the method of choice for a broad class of applications that involve complex multiphysics and multiscale phenomena. In fact, the two approaches have complementary strengths and weaknesses that can be exploited to compose efficient and robust strategies for solving nonlinear systems of equations. The expense of a Newton-Krylov method can be mitigated by using it as the coarse grid solver for FAS, which will often provide a good initial approximation. An improved coarse grid solver can also improve the

robustness of a nonlinear multigrid scheme. Linear multigrid methods can be used as preconditioners for a Newton-Krylov scheme, whether it is used as a coarse grid solver or as a standalone solver. As a standalone solver, Newton-Krylov schemes with multigrid preconditioning can make the choice of a smoother less critical. Further, no explicit linearization is needed to make these strategies effective.

Even within this framework, numerous questions still need to be addressed. What is the best multigrid cycling strategy to use as a preconditioner? How accurately should the FAS coarse grid problem be solved? Where in the grid hierarchy is the Newton-Krylov method most effective? What software components are needed, and how should they be organized, to facilitate exploration of these issues? Few analytic results are available to point to the most effective strategy, but numerical experiments help to identify fruitful avenues to investigate.

September 28, 1998

# Parallel PIC Modeling of Semiclassical Quantum Models

*Dean Dauger*

*University of California, Los Angeles  
dauger@physics.ucla.edu*

## Abstract

A new simulation model based on the semiclassical approach has been constructed to investigate the dynamics of quantum mechanical systems. The simulation code, which is under development, uses classical particle computations to approximate Feynman path integrals that evolve quantum mechanical wavefunctions forward in time. Interactions are treated semiclassically. The algorithm for the evolution of the classical particle trajectories is based on a plasma particle-in-cell code that is designed for parallel computers and is scalable to many processors. This will facilitate the modeling of many particle (100–1000) quantum systems. Progress will be reported on one-dimensional single and multiparticle systems and comparisons with physical situations with known analytic solutions will also be discussed.

September 23, 1998

# Fermion Monte Carlo

*Mal Kalos**Cornell University  
Theory & Simulation Science & Engineering Center*

## Abstract

We review the fundamental challenge of fermion Monte Carlo for continuous systems, the "sign problem," and some of the proposals that have been made for its solution, including some approximate schemes and others whose computing requirements grow exponentially. The issue is to find not the fundamental eigenmode of the Schrödinger equation in many dimensions, but one that has a special inversion of sign—antisymmetry in the exchange of pairs of coordinates. Thus the function sought is not everywhere positive.

We also describe a class of methods that depend upon the use of correlated dynamics for ensembles of walkers that carry opposite signs. We discuss the algorithmic symmetry between such walkers that must be broken to create a method that is both exact and as effective as for treating the fundamental symmetric solution. We explain the concept of marginally correct dynamics. Stable overlaps with an antisymmetric trial function given by such dynamics correspond to the lowest antisymmetric mode. Many-body harmonic oscillator problems are particularly tractable: Their stochastic dynamics permits the use of regular geometric structures for the ensembles, structures that are stable when appropriate correlations are introduced, and that avoid the decay of signal-to-noise that is a normal characteristic of the sign problem. Finally, we outline a new generalization of the method for arbitrary potentials and describe the success in treating few-electron systems, free fermions and small systems of  $^3\text{H}$ .

September 17, 1998

# Developing Effective Communication Abstractions for Clusters of SMPs

*Steven Lumetta*

*University of Illinois at Urbana-Champaign  
steve@crhc.uiuc.edu*

## Abstract

Clusters of SMPs, or Clumps, present a two-level hierarchy for interprocess communication. Effectively addressing this hierarchy requires a substantial effort. Clearly, many applications will require that a programmer recognize the hierarchy to some extent to obtain optimal performance, but an effective abstraction can relieve the programmer of the bulk of the workload.

In the first part of my talk, I present a lightweight message-passing layer that abstracts communication into a single interface, transparently directing communication traffic through the appropriate medium. I describe the design, implementation, and performance of this multi-protocol layer on a cluster of Sun Enterprise 5000 servers connected by Myrinet. Applications within an SMP benefit from the use of fast communication, but resource contention can dominate performance on a Clump. In the second half of my talk, I present models of application performance and argue that realizing the full potential of the Clumps architecture may require that we abandon the global phase structuring that has proven so effective in developing parallel applications.



September 11, 1998

# PLUM: Parallel Load Balancing for Adaptive Unstructured Meshes

*Leonid Oliker*

NASA Ames Research Center  
oliker@riacs.edu

## Abstract

Dynamic mesh adaption on unstructured grids is a powerful tool for computing unsteady three-dimensional problems that require grid modifications to efficiently resolve solution features. By locally refining and coarsening the mesh to capture flowfield phenomena of interest, such procedures make standard computational methods more cost effective. Unfortunately, the adaptive solution of unsteady problems causes load imbalance among processors on a parallel machine. An efficient parallel implementation of such methods is extremely difficult to achieve, primarily because of the dynamically changing, nonuniform grid.

This talk will present the development of PLUM, an automatic portable framework for performing adaptive numerical computations in a message-passing environment. Mesh adaption, repartitioning, processor assignment, and remapping are critical components of the framework that must be accomplished rapidly and efficiently so as not to cause a significant overhead to the numerical simulation. PLUM requires that all data be globally redistributed after each mesh adaption to achieve load balance. An algorithm is presented for minimizing this remapping overhead by guaranteeing an optimal processor assignment. Results demonstrate that applying a processor reassignment algorithm to the default mapping of the parallel partitioner can significantly reduce the data redistribution cost.

A data redistribution model is also presented that predicts the remapping cost on the SP2. This model is required to

determine whether the gain from a balanced workload distribution offsets the cost of data movement. Portability is examined by comparing performance on a SP2, and Origin2000, and a T3E. A CCNUMA implementation of this work is currently underway. Preliminary shared memory results on the Origin2000 will also be presented.

September 3, 1998

# GPFS - a Scaleable Parallel File System for the RS/6000 SP

*Roger Haskin*

*IBM Almaden Research Center*

## Abstract

General Parallel File System (GPFS) is a scaleable parallel file system for the RS/6000 SP. GPFS allows large numbers of disks attached to multiple storage servers to be configured as a single file system. Files are striped across all disks in a file system for high throughput.

Programs running on multiple nodes in the SP can access data in a GPFS file in parallel. In addition to high-speed parallel file access, GPFS provides fault-tolerance. Working in conjunction with the SP2 Phoenix software, GPFS recovers automatically from disk and node failures. Its robust design makes GPFS appropriate for commercial applications such as large Web servers, data mining, and digital libraries. This talk presents the high-level architecture of GPFS, describes some of its advantages, and discusses plans for future work.

August 31, 1998

# Rayleigh-Taylor Mixing: Experiments and Simulations

*Yuan-Nan Young  
and  
Jonathan Dursi*

*University of Chicago*

## Abstract

Rayleigh-Taylor mixing, which occurs between unstably stratified fluids, is important in such areas as supernovae type Ia and inertial confinement fusion (ICF). Details of the mixing, however, are still poorly understood. We investigate the detailed spatial structure and size of the mixing zone by analyzing data from previous experiments on the Linear Electric Motor (LEM) using immiscible fluids, and results from numerical simulations with incompressible, Boussinesq fluid (low Atwood numbers) using a 2D pseudospectral code. We then compare differences and similarities between numerical and experimental results.

August 25, 1998

# X-Ray Diffraction and the Holographic Inverse Problem

*Abraham Szoke*

*LLNL Researcher Emeritus  
szoke1@llnl.gov*

## Abstract

X-ray diffraction on crystals is the most important tool for finding the three- dimensional structure of large molecules—proteins and DNA. There is a close parallel between x-ray diffraction and holography. Utilizing this parallel we have written a computer code that solves structures of large molecules. In this talk I will describe x-ray crystallography, holography, inverse problems and protein structure in one easy lesson.

August 24, 1998

# Algorithmic Scalability Issues in Computational Fluid Dynamics

*Timothy Barth*

NASA Ames Research Center  
barth@nas.nasa.gov

## Abstract

The first half of the presentation will focus on theoretical techniques for analyzing the performance of iterative methods for the CFD equations. In the second half of the presentation, I will discuss some ongoing work in parallel preconditioning for non-selfadjoint partial differential equations (PDEs).

Despite early success in the design of optimal complexity algorithms for elliptic PDE problems using multigrid (MG) and domain decomposition (DD) techniques, the design of optimal complexity algorithms in computational fluid dynamics (CFD) remains largely an open problem. Even among theoretically optimal elliptic PDE algorithms, computer architecture/hardware dependent characteristics such as memory latency and bandwidth conspire to prevent true parallel scalability from being achieved. The situation is more complex for the CFD equations where action at long distances takes place along characteristic directions.

To understand the performance of iterative methods for solving the discretized CFD equation, it is desirable to obtain energy and condition number properties for the symmetric portion of the discrete operator. In the present work, I consider systems of nonlinear conservation laws with convex extension together with stabilized finite element approximation. A general space-time discontinuous function space is assumed. Using nonlinear techniques together with specialized stabilization

techniques, it is possible to derive exact energy properties of these stabilized numerical methods for nonlinear conservation laws. Crucial to the analysis is the precise form of stabilization used in the finite element formulations. Using the analysis, subdomain and interface energies are then readily obtained from the theory. The exact energy balance equation is very revealing. For example, when applied to the nonlinear compressible Navier-Stokes equations, the exact analysis clearly identifies the ill-posed nature of the system energy for recirculating flow at high Reynolds number. The analysis also shows subtle deficiencies in the model problem analysis of Yavneh and Brandt for recirculating flow.

In the second portion of the presentation, I will consider preconditioning methods for convection-dominated fluid flow problems based on a nonoverlapping Schur complement domain decomposition procedure for arbitrary triangulated domains. This is joint work with Tony Chan (UCLA) and W.-P. Tang (University of Waterloo). The triangulation is first partitioned into a number of subdomains and interfaces, which induces a natural  $2 \times 2$  partitioning of the PDE discretization matrix. We view the Schur complement induced by this partitioning as an algebraically derived coarse space approximation. This avoids the known difficulties associated with the direct formation of an effective coarse discretization for advection dominated equations. By considering various approximations of the block factorization of the  $2 \times 2$  system, we have developed a family of robust preconditioning techniques. These approximations are introduced to improve both the sequential and parallel efficiency of the method without significantly degrading the quality of the preconditioner. The specific approximations that we have used include ILU-preconditioned GMRES subdomain solves, localized approximation of the interface Schur complement, and limited level-fill ILU interface backsolves. A number of 2D CFD calculations will be presented for both scalar advection-diffusion equations and the Euler equations.

August 21, 1998

# The Trials and Tribulations of Scriptable Scientific Software

*David Beazley*

*University of Chicago*  
*beazley@cs.cs.utah.edu*

## Abstract

Scripting languages such as Perl, Python, and Tcl are becoming an increasingly popular tool for the creation of flexible scientific software. Much of this popularity is due to the fact that scripting languages provide scientists with an interpreted environment that can be used for exploratory problem solving similar to that found in packages such as MATLAB, Mathematica, Maple, and IDL. In addition, scripting languages often simplify software development because they encourage the use of software components and provide a high-level environment for debugging and testing.

In this talk, I describe the use of scripting languages with the SPaSM large-scale molecular dynamics code at Los Alamos National Laboratory. Originally developed for the Connection Machine 5, SPaSM initially proved to be too difficult to use and maintain to be of practical value to scientists. However, the use of scripting languages effectively transformed this application into a highly flexible system that is now being used on a daily basis. Various aspects of the scripting environment will be described including an integrated data analysis and visualization component, remote simulation monitoring over the Internet, and the use of scripting languages on parallel machines such as the Avalon DEC-Alpha Linux cluster at Los Alamos (a price-performance finalist in the 1998 Gordon-Bell Prize).

Currently, SPaSM utilizes Python ([www.python.org](http://www.python.org)) for its scripting interface. In addition, the SWIG interface generator ([www.swig.org](http://www.swig.org)) is used to construct Python interfaces to C/C++ libraries. Much of the talk will focus on the impact of using these tools over a three-year period. In addition, limitations and future challenges will be discussed.

August 19, 1998

# The SCIRun Problem Solving Environment

*Steve Parker*

*University of Utah ASCL Alliance Center  
sparker@taz.cs.utah.edu*

## Abstract

Computational steering has been defined as “the capacity to control the execution of long-running, resource-intensive programs.”

In the field of computational science, we apply this concept to link visualization with computation and geometric design to interactively explore (steer) a simulation in time and/or space. As knowledge is gained, a scientist can change the input conditions and/or other parameters of the simulation.

Although steering was proposed over a decade ago, it is only gradually becoming a popular paradigm for scientific computing. Computational steering is difficult because it requires in-depth knowledge in a wide range of disciplines from geometric modeling to scientific computing to scientific visualization and graphics. Most scientists do not have the necessary expertise in visualization, and most visualization experts do not perform large-scale scientific simulations. In order to successfully apply computational steering to these iterative design problems, we implement a problem solving environment (PSE), called SCIRun, wherein these various phases of the scientific computing process may be integrated.

Implementation of a computational steering framework requires a successful integration of the many aspects of scientific computing, including geometric modeling, numerical analysis, and scientific visualization. All aspects must be effectively coordinated within an

efficient computing environment (which, for large-scale problems, means dealing with the subtleties of various high-performance architectures).

In this talk, I will describe how the architecture of SCIRun addresses these problems.

August 14, 1998

# Scaling Heterogeneous Information Access for Wide-Area Environments

*Louisa Raschid*

*University of Maryland*  
*louisa@umiacs.umd.edu*

and

*Michael Franklin*

*University of Maryland*  
*franklin@cs.umd.edu*

## Abstract

Research in data integration technology aims to enable seamless access to data stored in a wide variety of repositories. Recently, numerous successful prototype systems have been developed. As such systems begin to be deployed in a wide-area network-based environment, however, they will encounter significant challenges arising from the huge number of disparate, unpredictable, and unreliable repositories.

In order for the deployment of such systems to be successful, data integration technology must overcome many scalability problems, including

- The capabilities and contents of heterogeneous repositories can be very dissimilar, making it difficult to generate working access plans.
- The availability and response time for accessing remote repositories can fluctuate dramatically.
- Little support currently exists for identifying and locating repositories that are relevant to a particular application.

In this talk, we describe techniques that we have developed for addressing these scalability problems, including

- Tools for generating wrappers, specifying source capabilities, and generating alternative access plans.
- Query Scrambling, a reactive approach to query execution that can adapt to changes in the environment.
- WebSemantics, an architecture for publishing, locating, and transparent access to sources via the WWW. We then discuss our on-going research on metadata management for the discovery, querying, and dissemination of data sources on the Internet.



August 12, 1998

# Real Gas Effects in Oblique Shock Wave Reflection Experiments and Computations

*Harland Glaz*

*University of Maryland Department of Mathematics  
hmg@maria.umd.edu*

## Abstract

An extensive series of shock wave reflection experiments was performed (1970–1990, approximately) at the University of Toronto Institute for Aerospace Studies (UTIAS) under the direction of Professor I. I. Glass. Test gases included air, CO<sub>2</sub>, and SF<sub>6</sub>; whole flowfield data is available from interferograms. Since 1975, CFD simulations of these experiments have been widely undertaken for various purposes, and with substantial success.

This talk will focus on a series of “anomalous” experiments in heavy test gases (such as SF<sub>6</sub> and isobutane) run at high shock-wave Mach number, very low ambient test gas density and pressure, and at or near 37° wedge angle. Interferograms and computational results using second-order Godunov schemes will be presented and compared. It will be seen that much of the phenomenology is still poorly understood. Computational results attempting to partially resolve the situation via inviscid simulations, but with a nonconvex EOS, are presented, including new results by Bei Wang ('98 Ph.D. dissertation, UMD). Other, more speculative, ideas will be discussed as well.

August 7, 1998

# Experiment Management Support for Parallel Performance Tuning

*Karen L. Karavanic*

*University of Wisconsin, Madison  
karavan@cs.wisc.edu*

## Abstract

Existing performance tools focus on gathering and reporting information about a single execution of a program. However, the development of a high-performance parallel system or application is an evolutionary process: Both the code and the environment go through many changes during a program's lifetime. At each change, a key question for developers is, how and how much did the performance change?

My research reframes performance tuning as a specialized instance of scientific experimentation, and develops methods for storing, viewing, and using performance data that span a variety of executions, program versions, and environments. In this talk I shall report some early results with a performance tuning study and with a scientific application run in changing environments. I shall also describe work in progress that investigates the use of historical performance data to improve automated performance diagnosis.

July 28, 1998

# Preconditioning for the Steady-State Navier–Stokes Equations with Low Viscosity

*Howard Elman*

*University of Maryland  
elman@cs.umd.edu*

## Abstract

We present a preconditioner for the linearized Navier–Stokes equations that is effective when either the discretization mesh size or the viscosity approaches zero. For constant coefficient problems with periodic boundary conditions, we show that the preconditioning yields a system with a single eigenvalue equal to one, so that performance is independent of both viscosity and mesh size.

For other boundary conditions, we demonstrate empirically that convergence depends only mildly on these parameters and we give a partial analysis of this phenomenon. We also show that some expensive subsidiary computations required by the new method can be replaced by inexpensive approximate versions of these tasks based on iteration, with virtually no degradation of performance.

July 22, 1998

# Model-Based Object Recognition Using Algebraic Functions of Views

*George Bebis*

*University of Nevada, Reno*

## Abstract

The recognition and classification of objects is a spontaneous, natural activity for many biological vision systems. In contrast, building systems capable of recognizing relevant 3D objects in their environment, with accuracy and robustness, has been a difficult and challenging task in computer vision. Object recognition is difficult because the appearance of an object can have a large range of variation due to photometric effects, scene clutter, changes in shape (e.g., the object is not rigid) but most importantly, due to viewpoint changes. This results in numerous different images even for the same object.

Accommodating variations due to viewpoint changes is a central problem in the design of any object recognition system, and one whose solution is likely to have implications throughout the system. In this talk, we will present a new technique for object recognition based on the recently proposed theory of "algebraic functions of views," which provides a powerful mathematical foundation for tackling variations in the appearance of a 3D object's shape due to viewpoint changes.

July 17, 1998

# Synthesis of Linear Phase Multirate Filter Banks for Signal and Image Coding Applications

*Chris Brislawn*

*Los Alamos National Laboratory  
brislawn@lanl.gov*

## Abstract

Wavelet transforms and their digital signal processing counterparts, multirate filter banks, have proven superior to block transform coding for a variety of digital data coding and are now being incorporated into a number of commercial and international standards for audio, still and moving picture data compression. The speaker was involved in drafting the first such wavelet-based standard, the FBI's national specification for coding digitized fingerprint images.

This talk will review the definitions of wavelet transforms and multirate filter banks and indicate how they are used in digital signal and image coding. We will then present a new cascade-form architecture for synthesizing linear phase filter banks of one of the classes commonly used in image compression applications. The new architecture, which improves on earlier work of Vaidyanathan and Nguyen, achieves a reduction of 33% to 50% (asymptotically) in the number of multiplications per unit input needed to implement the filter bank when compared to direct form implementation.

July 10, 1998

# Legion – An Applications Perspective

*Andrew Grimshaw*

*University of Virginia  
grimshaw@cs.virginia.edu*

## Abstract

Legion is a reflective metasystem project at the University of Virginia. It is designed to provide users with a transparent interface to resources in a wide-area system, both at the programming interface level and at the user level. Legion addresses issues such as parallelism, fault-tolerance, security, autonomy, heterogeneity, resource management, and access transparency in a multi-language environment. While fully supporting existing codes written in MPI and PVM, Legion provides features and services that allow users to take advantage of much larger, more complex resource pools.

With Legion, for example, a user can run a computation on a supercomputer at a national center while dynamically visualizing the results on a local machine. As another example, Legion makes it trivial to schedule and run a large parameter space study on several workstation farms simultaneously. Legion permits computational scientists to use cycles wherever they are, allowing larger jobs to run in shorter times through higher degrees of parallelization.

July 1, 1998

# Hierarchical Approaches for Representing and Visualizing Massive Scientific Data Sets

*Bernd Hamann*

*University of California Davis  
Department of Computer Science*

*Co-Director of the Center for Image Processing and  
Integrated Computing (CIPIC)  
hamann@cs.ucdavis.edu*

## Abstract

One of the most challenging and important problems that the science and engineering communities are facing today—and even more so in the future—is representing, visualizing, and interpreting very large data sets. Such data sets commonly result from computer simulations of complex physical phenomena (e.g., computational physics, climate modeling, or ocean modeling) and from high-resolution imaging (e.g., satellite imaging or medical imaging).

The technology currently being used to represent massive data sets is inappropriate for interactive and efficient data analysis and visualization. It is impossible for a user of a visualization system to "navigate" through a data set consisting of several million (or billion) data points, and analyze the data set entirely. In this talk, I will present various ideas that seem promising in the context of overcoming some of the problems associated with the visualization of very large data sets. I will point out the necessity of bringing together ideas from approximation theory and geometric modeling (splines),

computational geometry (tessellations), optimization (simulated annealing), and other related fields. I will point out avenues for representing massive data sets using hierarchical approaches that facilitate visualization and analysis.

June 29, 1998

# Least Squares Finite Element Approximations Based on Minus One Inner Product, and Their Analysis

*Raytcho Lazarov*

*Texas A & M University  
lazarov@math.tamu.edu*

*A presentation of joint work conducted with  
J. Bramble and J. Pasciak*

## Abstract

The merits and the deficiencies of various least squares formulations for solving boundary value problems will be discussed as a short introduction. Major advantages of the least squares methods are the stability, symmetry and positive definiteness of the corresponding discretizations.

The main objective of this talk is to introduce and to analyze least squares formulations for second order elliptic problems based on a discrete minus one inner product. The proposed formulation does not require the classical condition of Ladyzhenskaya–Babuska–Brezzi. Using these functionals, we develop finite element approximations for non-symmetric and indefinite second order elliptic problems, and for the equations of incompressible and almost incompressible elastic media. The method leads to symmetric and positive definite algebraic problems of optimal rate of convergence.

This approach is based on the recent theoretical advances and practical implementations of multilevel and multigrid methods, multilevel splittings of Sobolev norms, and domain decomposition methods.



June 25, 1998

# Efficient Elliptic Solvers for Supercomputers Based on High- Performance Microprocessors

*Ulrich Ruede*

*Institut fuer Mathematik,  
Universitaet Augsburg, D-86135 Augsburg, Germany  
Ulrich. Ruede@Math.Uni-Augsburg.DE*

## Abstract

We recognize the trend to build even supercomputers using standard microprocessors. Such processors are amazingly fast in terms of computational power. The performance of such systems, however, is often limited by the speed of the memory system. This is a severe bottleneck for all conventionally programmed iterative methods, since they exhibit only limited re-use of data from faster cache memory.

Computational experiments can show the performance of different data structures and algorithmic patterns for current microprocessors and their memory systems. To increase the efficiency, the design of iterative methods must avoid patterns that limit on-chip parallelism, and the algorithms must be restructured such that the re-use of cached data is improved. With a carefully designed multigrid method, a system with one million unknowns can be solved in approximately one second on a single CPU node.

June 19, 1998

# High-Performance Thread Libraries: Past and Current Work

*Dirk Grunwald*

*University of Colorado Boulder  
grunwald@foobar.cs.colorado.edu*

## Abstract

Computer runtime systems use thread libraries to provide a way for programmers to express concurrency. On sequential processors, a thread library can be used to mask communication latency by switching threads on communication stalls. On shared memory parallel processors, threads are used to provide work for different processors. In this talk, I'll describe two studies to improve the performance of thread systems.

The first study<sup>1</sup> improves the mechanics of thread libraries, particularly for fine-grained threads. For conventional languages, threads require space for the per-thread stack and must save registers across context-switch boundaries. We used whole-program optimization to reduce the needed space, and, as a side effect, the number of TLB misses encountered by threaded programs. We also used link-time inter-procedural live-register analysis to determine and reduce the number of live registers that need to be preserved around a voluntary context switch. Both of these mechanisms were implemented on a DEC Alpha system using the "OM" binary optimization system from DEC Western Research Labs, and speedups of 15–35% were demonstrated.

The second study<sup>2</sup> was designed to reduce the execution cost and simplify the scheduling of scientific applications on parallel systems. We designed a new thread abstraction, called Iterates, that represents the execution of a subspace of an iteration space. Iterates are linked by

dependence constraints that are resolved at execution time. This combination effectively turns conventional applications into "large-grain data flow" execution models, allowing the scheduling system to implement load balancing, affinity scheduling and automatic wave fronting. A prototype system was implemented on a DEC Alpha system, and achieved speedups of 10–40% on various application kernels.

In closing, I'll also mention an on-going project to build a high-performance distributed object system based on the CORBA distributed object model.

<sup>1</sup>Joint work with Richard Neves, now at the IBM T.J. Watson Research Center, Yorktown Heights, NY.

<sup>2</sup>Joint work with Suvas Vajracharya, now at Los Alamos National Laboratory.

June 18, 1998

# Reduced-Order Modeling for Multi-Input Multi-Output Linear Dynamical Systems

*Roland Freund*

*Bell Laboratories*  
*freund@research.bell-labs.com*

## Abstract

Multi-input multi-output time-invariant linear dynamical systems arise in important applications, for example, in the simulation of electronic circuits. The state-space dimensions of such systems can be so large that time-domain integration of the original system would be inefficient or even prohibitive. Instead, the original system is replaced by a suitable reduced-order model whose state-space dimension is small enough that it can be solved numerically.

In this talk, we discuss reduced-order modeling techniques based on matrix-Pade approximation of the frequency-domain transfer function of the multi-input multi-output linear dynamical system. We show how such reduced-order models can be computed in a stable and efficient way via a Lanczos-type method for multiple starting vectors. It is desirable and often crucial that reduced-order models inherit the essential properties, such as stability or passivity, of the original system. Next, we discuss stability and passivity of Pade-based reduced-order models. Finally, we present results of numerical experiments with linear dynamical systems arising in circuit simulation.

June 16, 1998

# High-Performance Computational Grids

*Ian Foster*

*Argonne National Laboratory  
and University of Chicago*

## Abstract

A computational grid, like its namesake the electric power grid, provides quasi-ubiquitous access to capabilities that cannot easily be replicated at network endpoints. In the case of a high-performance grid, these capabilities include both high-performance devices (networks, computers, storage devices, visualization devices, etc.) and unique services that depend on these devices, such as smart instruments, collaborative design spaces, and metacomputations.

In this talk, I discuss some of the technical challenges that arise when we attempt to build such grids: in particular, the frequent need to meet stringent end-to-end performance requirements despite a lack of global knowledge or control. I then introduce the Globus Project, a multi-institutional effort that is developing key grid infrastructure components, for authentication, resource location/allocation, process management, communication, and data management. I conclude with a discussion of our experiences developing test beds and applications based on Globus components.

June 12, 1998

# Database Middleware for Heterogeneous Data Sources

*Laura Myers Haas*

*IBM Almaden, Garlic Project  
LAURA@almaden.ibm.com*

## Abstract

Businesses today rely on large collections of data stored in diverse systems with differing capabilities. Database middleware systems can provide an integrated view of data stored on various relational systems, as well as a few non-relational sources. They do not help, however, when the data sources to be integrated include a broader range of sources, such as CAD/CAM systems, text search engines, molecular structure databases, or customer-specific repositories.

We describe Garlic, an IBM prototype that allows integration of diverse sources such as the above, and allows new sources to be easily added to an existing installation. Garlic offers the ability to interrelate data from multiple sources with a broad range of querying capabilities, in a single, cross-source query. A significant focus of the project is the provision of support for data sources that provide type-specific indexing and query capabilities, such as text search, or search by molecular structure. Garlic's "wrapper architecture" encapsulates data sources, allowing new sources to be added quickly, and accommodating a wide variety of sources, with a broad range of traditional and nontraditional query processing capabilities. Wrappers model legacy data as objects, participate in query planning, and provide standard interfaces for method invocation and query execution. Garlic provides an object-oriented query language (along the lines of SQL3), and extends database optimization technology to create efficient plans for queries over multiple sources, whatever their query capabilities, using wrapper input.

A Garlic prototype has been operational since 1995, and has been used as the basis of a customer joint study in the pharmaceutical domain. We are currently using Garlic as the basis of a new venture into scientific information management, initially for the petroleum industry. In this talk, we will describe the overall Garlic architecture, as well as our experiences to date when using Garlic.

June 11, 1998

# The Development of an Agglomeration Multigrid Method for Highly Anisotropic Unstructured Mesh CFD Problems

*Dimitri Mavriplis*

*Research Fellow  
ICASE, NASA Langley Research Center  
dimitri@icase.edu*

## Abstract

Agglomeration multigrid methods represent a variant of algebraic multigrid methods that can be applied directly to non-linear problems. This enables the solution of computational fluid dynamics (CFD) problems without the explicit storage of a Jacobian, which leads to large memory savings over the traditional approach of applying algebraic multigrid to the linearized system of equations. Similar to algebraic multigrid methods, agglomeration multigrid methods consist of a setup phase (coarsening) and a solution phase.

For inviscid flow problems where the underlying grid is isotropic, coarsening algorithms based on unweighted graphs have been shown to result in good overall convergence rates. For high-Reynolds number viscous flows, where extreme grid stretching is required to resolve the thin boundary layer and wake regions, directional coarsening is required, which can be implemented using a weighted graph algorithm. However, directional coarsening results in higher complexity coarse level meshes, which in turn requires additional memory and CPU time within a multigrid cycle. To overcome this problem, aggressive coarsening strategies have been developed, where coarse levels of

prescribed complexity reduction are generated. To maintain effective convergence rates, this strategy is coupled with a locally implicit multigrid smoother, which operates on locally constructed lines in the unstructured mesh, using a weighted graph algorithm.

Results will be presented that demonstrate fast convergence rates for aerodynamic Navier-Stokes flows that are independent of the grid stretching. In addition, the parallelization of this solver using the MPI message-passing library will be described, and scalability results on the CRAY T3E and SGI Origin 2000 will be given. Finally, the solution of a large-scale problem involving over 13 million grid points on the T3E will be described.

May 29, 1998

# Three Modest Proposals in Aid of Scaling PDE Solvers to Teraflops: Asynchronicity, Tensoricity, and Memtropy

*David E. Keyes*

*Old Dominion University  
Computer Science Department, and ICASE, NASA  
Langley Research Center  
keyes@cs.odu.edu*

## Abstract

Asynchronous iterative methods have somewhat of a “bad name” in the theoretical parallel algorithms literature, where worst case estimates for linear problems are pessimistic. Asynchronous methods, however, may deserve renewed consideration for heterogeneous nonlinear problems at ASCI scales. Heterogeneities (e.g., shocks, flame fronts, plastic zones) lead to “nonlinear stiffness,” (widely varying magnitudes in the nonlinearity tensor, the term after the Jacobian in the multivariate Taylor expansion upon which Newton's method is based).

Many important problems are strongly nonlinear in a small region that is embedded in an ambient region that is only weakly nonlinear. Treating heterogeneous problems bulk synchronously, as with the speaker's favorite Newton-Krylov-Schwarz methods, at high processor granularities either wastes cycles in the near-linear regions or enslaves load balancing to the evolving physics. Alternatively, it should be relatively easy to accommodate asynchronous forms of nonlinear

Schwarz iteration through multithreading, the natural software model for interactive visualization and computational steering. Each process “publishes” updates of solution values for on-demand inspection by processes governing neighboring subdomains, not by a global hand-shaking exchange.

Exploiting the possibilities of nonlinear Schwarz in an automated way requires a pointwise metric, analogous to a discretization error estimate in an adaptive discretization code, which quantifies the degree of nonlinearity. We propose a metric, “tensoricity,” for this purpose. Orthogonally, but motivated by the same challenge of developing PDE solvers effective at Teraflops scales, we propose a new metric, “memtropy,” to provide architecture-independent rankings of data locality in programs destined for deep-memory hierarchy machines. No architecture-independent metric can be expected to be a reliable performance predictor on a given machine, but memtropy should be convenient in assessing trends when laying out data for non-flat memory systems, generally.

No performance data justifying the approaches advocated herein will be given yet; however, the necessity of moving beyond synchronous SPMD implementations of PDE solvers will be argued from performance data from large-scale aerodynamics runs. Moreover, the advocated looser approaches potentially lead to better machine utilization in the multi-physics computations ultimately required in ASCI.

May 22, 1998

# A Subspace-Based Model for Information Retrieval with Applications in Latent Semantic Indexing

*Hongyan Zha*

*Pennsylvania State University*

## Abstract

A theoretical foundation for latent semantic indexing (LSI) is proposed by adapting a model first used in array signal processing to the context of information retrieval using the concept of subspaces. It is shown that this subspace-based model, when coupled with the minimal description length (MDL) principle, leads to a statistical test to determine the dimensions of the latent-concept subspaces in LSI.

The effect of weighting on the choice of the optimal dimensions of latent-concept subspaces is illustrated. It is also shown that the model imposes a so-called “low-rank-plus-shift structure” that is approximately satisfied by the cross-product of the term-document matrices. This structure can be exploited to give a more accurate updating scheme for LSI, and to correct some of the misconception about the achievable retrieval accuracy in LSI updating. It is further demonstrated that based on the low-rank-plus-shift structure a divide-and-conquer method can be devised to compute the partial singular value decomposition (SVD) of a large sparse term-document matrix. Possible extensions of the model to improve retrieval accuracy are also pointed out.



May 19, 1998

# Problem-Solving Environments and NetSolve: A Network Server for Solving Computational Science Problems

*Jack Dongarra*

*University of Tennessee*

## Abstract

This talk presents a system called NetSolve that allows users to access computational resources, such as hardware and software, distributed across the network. This project has been motivated by the need for an easy-to-use, efficient mechanism for using computational resources remotely. Ease of use is obtained as a result of different interfaces, some of which do not require any programming effort from the user.

Good performance is ensured by a load-balancing policy that enables NetSolve to use the computational resource available as efficiently as possible. NetSolve offers the ability to look for computational resources on a network, choose the best one available, solve a problem (with retry for fault-tolerance) and return the answer to the user.

May 1, 1998

# Project VisualEyes: Integrated Data Synthesis, Analysis and Visualization of Parallel Adaptive Simulations

*Chandrajit L. Bajaj*

*University of Texas at Austin*

## Abstract

Even with the use of parallel supercomputers, conventional approaches to simulation and visualization are weakly suited to terascale problem sizes, where discretized domain models extracted from large images are much too large for interactive analysis, visualization and where physical simulations could take days to weeks. One experiment is to extract and employ progressively encoded, hierarchical meshes to support both adaptive simulation and interrogative visualization.

The VisualEyes project is focused on research in some of the core technologies (e.g., compression, adaptive meshing, interrogative visualization, and error estimation) as well as in integrating parallel simulation, with data analysis and collaborative visualization of multivariate scalar, vector, tensor fields. In this talk I shall present details of our progressive mesh encoding scheme, as well as the use of vector field topology for several applications in interrogative visualization. This project is driven by close cooperation with TICAM application developers in composite materials, electromagnetic scattering, and reservoir modeling.

April 28, 1998

# Coherent Structure Theory and Voting Rules

*Dean H. Judson*

*University of Nevada, Reno  
Statistics and Research Methods Laboratory  
and Social Psychology Ph.D. Program*

## Abstract

This presentation is a case study in systems isomorphy or systems analogy. I present the basic terms of coherent structure theory, which was originally developed for engineering reliability applications. I then show how the elements of coherent structures (components, minimal paths, modules, and monotonic structures) can be given a coalition or voting interpretation (decision-makers, minimal winning coalitions, subcommittees, or irrelevant group members).

I describe some basic theorems derivable from coherent structure axioms, and I discuss some implications of interpreting the theorems with respect to voting rules. Finally, I describe a recent project involving inferring the underlying coherent structure from a sample of data on components and system outcomes.

April 20 and 27, May 4 and 11, 1998

# Large-Scale Simulations in Fortran 95: An Object-Based Approach

*Paul F. Dubois*

*Lawrence Livermore National Laboratory  
X-Division  
dubois1@llnl.gov*

## Abstract

The course is intended for experienced Fortran 77 users. After this "free upgrade" to your skills, you should be able to use the full power of Fortran 95 and understand how to use an object-based approach to large-code architecture.

If you do not currently have any Fortran 90/95 reference materials, consider the following: *Fortran 90 Programming*, T.M.R. Ellis, Ivor R. Philips, and Thomas M. Lahey, Addison-Wesley, Reading, MA., 1994, ISBN 0-201-054446-6. If you search an on-line bookstore such as amazon.com with the word "Fortran" you'll find plenty more to choose from. If you have access to a Windows 95/NT machine, strongly consider the Digital Visual Fortran integrated environment. The book that comes with this is very good. (If you get one; one person reports that he didn't get anything except a CD.)

April 17, 1998

# Representation and Image Comparison Metrics for Large Data Sets

*Raghu Machiraju*

*Mississippi State University  
NSF Engineering Research Center on Computational  
Field Simulation*

## Abstract

In this talk I address two important operations needed for the visualization of large (terascale?) data sets, namely compressed domain representation, and metrics for comparing data sets and images. We first show how multiresolutional schemes based on wavelets can be used to represent large data sets.

Wavelets are used to detect coherent structures, which persist across scales. The sub-combining technique can be employed to identify regions in a volume that contain significant structures. These regions can be spatially partitioned (with octree) into blocks and each block can be coded again with a wavelet transform. The coefficients can now be packed into a bitstream in such a way that a ranking of perceptually significant structures is obtained. We show an example use of this representation for rendering iso-surfaces.

Using the same multiresolution representation scheme, we design image comparison metrics. We describe how the metrics respond to operators normally used in rendering. Finally, we show how these metrics can be extended to multi-part metrics that can be used to compare images and data sets alike and provide a better understanding of differences. Such metrics find use in navigating and browsing

April 17, 1998

# Change Detection in Data Warehousing and a Performance Study

*Nabil Rashad Adam*

*and*

*Igg Adiwijaya*

*Center for Information Management Integration  
and Connectivity (CIMIC),  
Rutgers University, Newark, NJ*

## Abstract

The change detection problem is encountered in various areas such as digital libraries and electronic commerce, specifically in data warehousing and Internet Web pages, among others. In such systems, the change detection problem manifests itself into two possible forms. First, there is a need to detect and propagate relevant changes to the underlying information sources to ensure data consistency. Second, such systems encounter continuous change not only in the content of Web pages but also to the Web pages themselves as new pages may be added and existing pages may be removed.

In this talk we address these two forms of change detection, i.e., detecting and propagating changes to relevant data (information sources in the former case and Web pages in the latter case).

April 14, 1998

# Multiresolution Algorithms in Computer Graphics

*Peter Schroeder*

*California Institute of Technology*

## Abstract

Computational techniques based on wavelets and more general multiresolution approaches have made tremendous inroads into computer graphics applications. Examples from modeling, simulation, and rendering include hierarchical surface editing, illumination computations, and level-of-detail display, among others. The success of these techniques is based on the favorable scaling properties of multiresolution transforms, their flexible time/accuracy tradeoffs, and recent advances in very general multiresolution construction techniques.

In this talk I will give an overview of the ideas behind these constructions, and use application examples from radiative transport, large -scale data compression, and arbitrary topology surface modeling, to illustrate their use effectiveness.

April 10, 1998

# Robust Simplification Methods for Triangle and Tetrahedral Meshes

*Kenneth I. Joy*

*University of California, Davis  
Computing Science Department Center for Image  
Processing and Integrated Computing*

## Abstract

One of the most critical and fundamental research problems encountered in the analysis and visualization of massive data sets is the development of methods for storing, approximating, and rendering large sets of data efficiently. The problem is to develop different representations of the data set, each of which can be substituted for the complete set depending on the requirements of the analysis or the visualization technique.

In this talk we present new methods for the construction of multiple levels of triangle and tetrahedral meshes. Starting with an initial, high-resolution triangulation, we construct coarser representation levels by collapsing simplices in the mesh. Based on weights defined for each simplex, we identify the simplex whose elimination would cause a minimal increase in error, and collapse this object. Weights are stored for individual simplices and are updated as the mesh is simplified. Different strategies can be used in the simplification process to produce hierarchies of meshes that can be used for different applications. These methods result in a hierarchical data description suited for the efficient visualization of large data sets at varying levels of detail.



March 30, 1998

# The Challenge of Fermion Monte Carlo

*Malvin H. Kalos*

*Cornell University  
Theory & Simulation Science & Engineering Center*

## Abstract

We will review the fundamental challenge of fermion Monte Carlo for continuous systems, the “sign problem,” and some of the proposals that have been made for its solution. In particular, we describe methods that depend upon the use of correlated dynamics for ensembles of correlated walkers that carry opposite signs. We discuss the algorithmic symmetry between such walkers that must be broken to create a method that is both exact and as effective as for bosonic systems. We explain the concept of marginally correct dynamics. Stable overlaps with an antisymmetric trial function given by such dynamics correspond to the correct fermion ground state.

Many-body harmonic oscillator problems are particularly tractable: Their stochastic dynamics permit the use of regular geometric structures for the ensembles. These structures are stable when appropriate correlations are introduced, and avoid the decay of signal-to-noise that is a normal characteristic of the sign problem.

Finally, we outline a generalization of the method for arbitrary potentials, and describe the progress in treating few electron systems.

March 27, 1998

# Computational Analysis in Orthopedic Design

*Ann Hollister*

*Louisiana State University Medical Center*

*and*

*Louise Focht*

*Avanta Orthopedics, Inc.*

## Abstract

Anne Hollister, orthopedic surgeon at Louisiana State Medical Center, and Louise Focht, co-founder of Avanta Orthopedics, Inc. (San Diego) will present clinical and industry perspectives on the applications of computational modeling in prosthetic implant design for human joint replacements. Traditional design processes involve CAD development but do not include extensive computational analysis. One result has been that design flaws often do not appear until after many prostheses have been implanted, and patients have suffered from premature failure of these joint replacements. Properly executed, computational analysis can predict likelihood of failure of joint replacements and, when incorporated early in the design process, can improve the design quality.

Hollister is a practicing surgeon and has designed several prosthetic joints. Avanta Orthopedics is a company specializing in replacements for small joints, such as those found in the hand, elbow, and foot.

March 20, 1998

# Multigrid Methods for Unstructured Grids and for Convection– Diffusion

*Jinchao Xu*

*Penn State University  
Mathematics Department  
xu@math.psu.edu*

## Abstract

After a general introduction of multigrid methodology, a number of techniques will be presented in this talk on how a finite element equation posed on an unstructured grid in two or three dimensions can be solved within optimal computational complexity by special multigrid methods such as nonnested multigrid methods, auxiliary space methods and agglomeration methods. If time allows, some new discretization and iterative techniques will be discussed on convection dominated problems. Both theoretical and numerical results will be reported.

March 19, 1998

# Solving Eigenvalues Problems

*Jesse L. Barlow*

*Pennsylvania State University  
Department of Computer Science and Engineering  
barlow@cse.psu.edu*

## Abstract

This talk gives a discussion of recent work in the solution of eigenvalue problems. Two recent classes of results are surveyed: accuracy and perturbation theory, and methods for solving large-scale eigenvalue problems. In the first area, we discuss new ways of characterizing accurate solution of the symmetric eigenvalue and singular value problems. Traditionally, the singular value problem has been viewed through the window the symmetric eigenvalue problem, but it turns that it is better to characterize the accurate solution of the symmetric eigenvalue problem through the window of the singular value decomposition. In the solution of large--scale eigenvalue problems, we look at some recent results on approximations generated by Lanczos methods.

March 9, 1998

# Revolutionizing Parallel Programming through Better Hardware

*John Feo*

*Tera Computer Company*

## Abstract

The commercial availability of MPP and clustered SMP systems has made very high peak performance systems readily available to scientists. However, the difficulty of programming for data locality and the limitations of medium- and coarse-grain parallelisms have limited the usefulness of these systems. The Tera MTA is a scalable, shared memory, MIMD computer that supports very fine-grain synchronization. It has no local memory and no caches. Each processor is 128 virtual processors feeding a single instruction pipeline. As long as one of the virtual processors can issue an instruction every cycle, the processor remains fully utilized. Thus, parallelism is the only limiting factor on the MTA. Data locality, granularity, and task scheduling for load balancing are non-issues. The MTA sustains high execution rates over a wide spectrum of applications, while reducing programming costs and encouraging parallel algorithm development.

In this talk, I will describe the overall architecture of the MTA and the current status of the SDSC machine. I will explain how one develops and optimizes code for the system using the ASCI benchmark Sweep3D and other applications as time permits. I will give some preliminary performance numbers that compare quite favorably with other HPC systems.

February 26, 1998

# Information Retrieval via Limited-Memory Matrix Methods

*Tamara G. Kolda*

*University of Tennessee  
Department of Computer Science and Oak Ridge  
National Laboratory (ORNL) Householder Fellow,  
Mathematical Sciences Section  
kolda@msr.epm.ornl.gov*

## Abstract

With ever-larger collections of documents available electronically, a need has arisen for fast and efficient search engines. Latent semantic indexing (LSI) approximates a matrix representing a document collection using the truncated singular-value decomposition (SVD); this allows automatic recognition of latent relationships between words and leads to a more efficient search engine. We propose replacing the SVD with what we call the semi-discrete decomposition (SDD). The resulting SDD-based LSI performs as well as the SVD-based method, requires substantially less storage, and processes queries faster. Furthermore, the SDD is easy to update when new documents are added to the collection.

February 24, 1998

# An Adaptive Projection Method for Modeling Low Mach Number Flows

*John B. Bell*

*Lawrence Berkeley National Laboratory (LBNL)  
Center for Computational Science and Engineering (CCSE)  
JBBell@lbl.gov,*

## Abstract

We describe an adaptive projection algorithm for low Mach number flows. The basic approach uses hierarchical grids that are refined in both space and time. In this presentation we will review the single grid projection methodology and discuss the key issues in the algorithm that must be addressed in developing an adaptive algorithm. With these issues in mind, we discuss the design principles that we use to solve partial differential equations on adaptive grids.

We then describe how these principles are used to develop an adaptive version of the projection algorithm. Numerical examples illustrating the capabilities of the method will be presented.

February 20, 1998

# Multi-Source Data Analysis in Science and Engineering

*Samuel P. Uselton*

*NASA Ames  
MRJ Technology Solutions  
uselson@nas.nasa.gov*

## Abstract

As digital data acquisition becomes easier, cheaper and more pervasive, and computational simulations gain increasing fidelity and detail, many activities can benefit from the combined analysis of data from several sources. Building a useful multi-source analysis system requires solving many problems, some of a pragmatic, engineering nature, and some of a more basic nature. The issues seem remarkably similar whether the application is weather modeling, environmental assessment and remediation planning, oil exploration and production, or engineering design processes. The problems to be addressed range from efficient access to large amounts of data from multiple heterogeneous sources, to design of user interfaces, and inventing visualization techniques. Integrated display and comparative analysis of relevant data is interesting and relatively unexplored.

Our multi-source visualization (MSV) project uses a specific problem domain, concurrent design of aircraft, to focus research and development efforts in this area. Frequent contacts with the Earth observation system (EOS) project, meso-scale atmospheric modeling researchers, and nanotechnology researchers, among others, keep broader needs in view.

All the work done by the Data Analysis Group of the NAS Division at NASA Ames Research Center is relevant and useful to this project. This group's work ranges from specific visualization techniques, through innovative user interfaces, the software engineering required to build

complete systems, systems level performance improvements, effective access to very large data sets, and exploitation of large heterogeneous collections of scientific and engineering data. Much of this work will be touched upon, but the focus will be on the work that is specifically driven by the need for a variety of users to exploit common collections of data from many diverse sources.

This work is very much in progress. There are results to show, but also problems not yet satisfactorily resolved. This talk will be profusely illustrated, of course.



February 20, 1998

# Simulating Turbulence and Turbulent Convection on SMP Clusters, with the PPM Gas Dynamics Code

*Paul Woodward*

*University of Minnesota  
Laboratory for Computational Science  
and Engineering (LCSE)*

## Abstract

The PPM gas dynamics code is being used in an extensive program of simulations of turbulent thermal convection in stars. During the last year it has been possible to treat entire 3D model stars, using this code on large SMP clusters. Using the ASCI machines at Los Alamos, detailed simulations of the compressible turbulence in its own right have also been carried out using the same techniques on grids of up to a billion cells. This work, the product of an interdisciplinary team at the University of Minnesota's Laboratory for Computational Science & Engineering collaborating with both Livermore and Los Alamos, gives us the opportunity to look inside stars and to observe the complex interactions of thermal convection with both rotation and pulsation of the stars.

The simulations of turbulence also offer unique views at unprecedented resolution of the dynamical processes at work. The methods used to make these calculations efficient on SMP clusters will be discussed and early visualizations of the results presented. These large simulations produce multi-terabyte data sets. Techniques used to deal with these data volumes will be discussed and directions of further research presented.

February 11, 1998

# Tecolote: An Object-Oriented Framework for Hydrodynamics

*Kathy Holian*

*Los Alamos National Laboratory  
Software Design and Development  
ksh@lanl.gov*

## Abstract

Tecolote is an object-oriented framework for both developing and accessing a variety of hydrodynamics models. It is written in C++, and is in turn built on the Parallel Object-Oriented Methods and Applications (POOMA) framework. The Tecolote framework is meant to provide modules (or building blocks) to put together hydrodynamics applications that can encompass a wide variety of physics models, numerical solution options, and underlying data storage schemes, although with only the necessary modules activated at runtime. Tecolote has been designed to separate physics from computer science, as much as humanly possible.

The POOMA framework provides fields in C++ to Tecolote that are analogous to Fortran 90-like arrays in the way that they are used, but that, in addition, have underlying load balancing, message passing, and a special scheme for compact data storage. The POOMA fields can also have unique meshes associated with them that can allow more options than just the normal, regularly spaced Cartesian mesh. They also permit one, two, or three dimensions to be immediately accessible to the code developer and code user.

Having a framework in which to develop hydrodynamics methods furnishes numerous advantages. The most obvious is that one can build upon models (both physical and numerical) that have already been developed and tested. It is also advantageous for the physics programmer to have complicated computer science issues already

solved, yet abstracted away. The types of computer problems that need to be solved these days, but that the physicist does not want to have to deal with, include such things as message passing and load balancing.

We have endeavored wherever possible to write the framework in as object-oriented a manner as possible so that model developers can enjoy maximal code reuse. This has also led to a nearly complete separation of the computer science from the physics coding. The framework has also been designed to be completely portable across a wide variety of platforms.

I will discuss the philosophy of the Tecolote framework, and will present a summary of the overall design, including a discussion of some of the objects that were used to put the framework together. The first hydrodynamics option using the framework that we are developing is a multi-material Eulerian code. This will allow us to compare with an existing hydrocode written in Fortran 90. I will present comparisons of both timings for large problems, and of possible maximum size of problem on the computer platforms of interest at LANL.

February 6, 1998

# Parallel Adaptive Mesh Refinement

*Dan Quinlan*

*Los Alamos National Laboratory  
Scientific Computing Group  
dquinlan@lanl.gov*

## Abstract

Adaptive mesh refinement is a numerical technique for locally tailoring the resolution of computational grids. AMR permits the addition of finer grids to the global computational grid in an adaptive way so as to permit locally more accurate computations or the removal of global error introduced by local singularities. AMR as a numerical technique, is largely independent of the equations being solved, though numerous numerical and algorithmic issues are involved and are the subject of significant research. Unfortunately, AMR is not common place due largely to its inherent complexity.

Adaptive mesh refinement computations are complicated by their dynamic nature. In the serial environment, they require substantial infrastructures to support the regridding processes, intergrid operations, and local bookkeeping of positions of grids relative to one another. In the parallel environment, the dynamic behavior is more problematic because it requires dynamic distribution support and load balancing. Parallel AMR is further complicated by the substantial task parallelism, in addition to the obvious data parallelism, this task parallelism requires additional infrastructure to support efficiently. The degree of parallelism is typically dependent upon the algorithms in use and the equations being solved. Different algorithms have significant compromises between computation and communication. Substantial research work is often required to define efficient methods and suitable infrastructure. The purpose of this talk is to introduce AMR++ as an object-oriented library which forms a part of the Overture framework, a much larger object-oriented numerical framework

developed and supported at Los Alamos National Laboratory and distributed on the Web for the last several years.

The parallel issues in AMR are the subject of special attention within AMR++. Overture provides simple mechanisms for the specification of the distribution of Overture grid data (via P++ distribution mechanisms). These distribution mechanisms are used by AMR++ and the multilevel load balancer (MLB) to define the distribution of the adaptive grid. Different AMR algorithms place different requirements upon the distribution of the adaptive grid in the parallel environment. AMR++ supports a couple of different distribution mechanisms specialized for the most common sorts of adaptive mesh refinement solution methods. The underlying mechanisms for defining new distributions, however, are readily available to the user as well. This allows for the tailoring of AMR++ parallel support for different sorts of applications or for research on distribution mechanisms more generally.

January 16, 1998

# Adaptive Finite Element Methods on Manifolds, with Applications in Elasticity and Relativity

*Michael Holst*

*University of California Irvine  
Department of Mathematics  
mholst@uci.edu*

## Abstract

We discuss the numerical treatment of coupled nonlinear elliptic systems on manifolds. Such systems arise for example, in elasticity models of biological membranes, and in general relativistic models of massive objects. We begin by reviewing some differential geometry, and by taking a brief look at these two applications. Weak formulations of covariant nonlinear elliptic systems on manifolds are then examined.

Finite element approximation theory on manifolds is then discussed, and a computer implementation called MC is described. MC is a dimension-independent, simplex-based, ANSI-C finite element code for the numerical treatment of covariant differential operators on  $d$ -manifolds ( $d = 2, 3, \dots$ ). MC also implements several of the features in the popular package PLTMG (designed for 2D problems in the plane), including a posteriori error estimation, adaptive simplex subdivision, global Newton methods, continuation, and multilevel methods. We describe some of the details of MC, and present some numerical examples for an elasticity problem, and for the elliptic constraints in the Einstein equations.

December 19, 1997

# Death and Taxes, Nets and Caches: Facing Inevitabilities in Parallel CFD Simulation

*David E. Keyes*

*Old Dominion University  
ICASE, NASA Langley Research Center  
keyes@cs.odu.edu*

## Abstract

Demands for massive memory and high speed typically accompany one another in scientific and engineering computations, linking space to time in algorithm design. Some degree of programmer control must be exerted over data layout in coding for scalable distributed memory machines (even when the memory is accessible through the model of a global shared address space).

Fortunately, the laws of nature often cooperate with a basic scaling law of computer architecture: The magnitude of interaction between two degrees of freedom in a physical system decays with their spatial separation; therefore, the frequency and volume of data exchange between different points in the computational domain can be allowed to decay with distance in a trade-off involving memory access overhead and the precision required in a final result (or the rate of convergence required from a preconditioner).

For model problems, this trade-off has been formalized in convergence theorems. We have been exploring it primarily experimentally, applying domain decomposition preconditioners to structured-grid and unstructured-grid problems in computational aerodynamics and acoustics, and keeping Amdahl's Law at bay through the benefits of cache locality. (See <http://www.cs.odu.edu/~keyes/nsf/>.)

December 16, 1997

# Challenges of Future High-End Computing

*David H. Bailey*

*NASA Ames Research Center*

## Abstract

The next major milestone in high-performance computing is a sustained rate of one Pflop/s (also written one petaflops, or 10<sup>15</sup> floating-point operations per second). In addition to prodigiously high computational performance, such systems must of necessity feature very large main memories, as well as comparably high I/O bandwidth, and huge mass storage facilities. The current consensus of scientists who have studied these issues is that "affordable" petaflops systems may be feasible by the year 2010, assuming that certain key technologies continue to progress at current rates.

One important question is whether applications can be structured to perform efficiently on such systems, which are expected to incorporate many thousands of processors and deeply hierarchical memory systems. To answer these questions, advanced performance modeling techniques, including simulation of future architectures and applications, may be required. It may also be necessary to formulate "latency tolerant algorithms" and other completely new algorithmic approaches for certain applications. This talk will give an overview of these challenges.

December 12, 1997

# Dynamic Distributed Arrays: An Infrastructure for Parallel Implementations of Adaptive Computational Algorithm

*James Browne*

*University of Texas at Austin  
Department of Computer Science*

## Abstract

A dynamic distributed array is an data abstraction that may be expanded or contracted in response to the requirements of an adaptive algorithm but which retains the operational semantics of a normal Fortran array even while distributed across multiple processors. This talk will define and describe two implementations of dynamic distributed arrays.

The first, the hierarchical dynamic distributed array or HDDA, is the basis of an infrastructure for parallel implementation of computations based on structured adaptive meshes. The second, scalable dynamic distributed array or SDDA, is the basis of an infrastructure for parallel implementations of computations based on unstructured adaptive meshes.

The HDDA and the SDDA are built on a common conceptual foundation. Use of the SDDA will be illustrated by formulation of an hp-adaptive finite element computation. Results reported are largely the work of Manish Parashar and Carter Edwards.

December 9, 1997

# Accurate Discretizations and Efficient Solvers for Heterogeneous Groundwater Flow Equations

*Thomas F. Russell*

*University of Colorado, Denver  
Department of Mathematics*

## Abstract

The ability to compute accurate velocities is important for applications of flow and transport codes. In highly heterogeneous porous media, this task is a difficult one for standard numerical methods. Irregular geological features, which suggest the use of irregular grids, and variable directions of anisotropy add to the challenge. Mixed finite element methods, which solve a continuity equation and a Darcy equation simultaneously for shape-function representations of pressure and velocity, can overcome these obstacles and produce accurate results.

The discrete formulations, however, are often complex, and the linear algebraic equations are not amenable to standard solvers and are difficult to solve as efficiently as equations from other methods. This has inhibited practical applications, especially in 3D. A control-volume variant of the lowest-order Raviart-Thomas mixed method is presented for general logically rectangular grids (2D quadrilaterals, 3D hexahedra). The Darcy equation is enforced on a cell-sized "tank" (control volume) around each degree of freedom of the velocity (edge in 2D, face in 3D). The

discrete equations are simple, easy to implement, and involve only cell pressures and edge or face fluxes (integrated normal velocities).

Numerical tests in 2D show second-order convergence of the edge fluxes whenever the exact solution is not singular. The accuracy far exceeds that of commonly used methods. Also presented is an efficient solver for the 3D equations. It uses a convenient basis for the divergence-free velocity shape functions to reduce the equations to a symmetric positive-definite system of smaller size than has been possible previously. This enables the 3D equations to be solved with effort comparable to that for other methods.



November 11, 1997

# The Hierarchical Basis Multigrid Algorithm and Incomplete LU Decomposition

*Randolph E. Bank*

*University of California, San Diego  
Department of Mathematics*

## Abstract

We present a new family of multigraph algorithms, ILU-MG, based upon an incomplete sparse matrix factorization using a particular ordering and allowing a limited amount of fill-in. While much of the motivation for multigraph comes from multigrid ideas, ILU-MG is not an typical algebraic multilevel method. The graph of the sparse matrix  $A$  is recursively coarsened by eliminating vertices using a graph model similar to Gaussian elimination. Incomplete factorizations are obtained by allowing only the fill-in generated by the vertex parents associated with each vertex.

Because they are graph based, in principle multigraph methods can be applied to general sparse matrices, and in particular do not require a grid hierarchy. We will present some numerical examples that show the behavior of the method is similar to the classical HBMG iteration.

November 10, 1997

# Multigrid and ILU Methods for Large Scale Computing

*Sing-Lok (Justin) Wan*

*University of California, Los Angeles  
Department of Mathematics*

## Abstract

This talk consists of two parts. In the first, we present a robust interpolation for multigrid methods based on the concepts of energy minimization and approximation. The formulation is general; it can be applied to any number of dimensions and to general computational domains. In 1D, we prove that the convergence rate of the resulting multigrid method is independent of the PDE coefficient and of the mesh size. In 2D, we demonstrate numerically the effectiveness of the multigrid method by applying it to rough coefficient problems and an unstructured grid problem.

In the second part, we present a parallelization of the ILU algorithm. The idea is based on multicoloring. We describe a parallel coloring algorithm by Jones and Plassmann and show how it can be applied to constructing incomplete LU factors for ILU(0). We also discuss an extension to ILU(1) and the issue of the memory allocation.



## Sabbatical Visitor

---





# Numerical Methods for PDEs, Large Scale Computer Simulations, and Applied Scientific Computations

*Raytcho Lazarov*

*Texas A&M University*

## Abstract

The researcher shall provide research assistance with projects at Lawrence Livermore National Laboratory (LLNL) that are related to numerical methods for PDEs, large scale computer simulations and applied scientific computations. The Center for Applied Scientific Computing (CASC) at LLNL has a strong and diverse program in design, testing and application of numerical algorithms in science and engineering.

The main direction of the research assistance will be with the following current projects at CASC

- Scaleable Linear Solvers.
- FOSLS (first order systems least squares).
- Modeling Subsurface Flow and Chemical Migration on High Performance Computers.
- Structured Adaptive Mesh Refinement Application.

Specifically, the researcher shall work on the following two topics.

### *Finite volume element methods for convection–diffusion–reaction problems.*


This will include developing and studying various approximation strategies in finite volume methods as locally conservative schemes and also design, testing and implementation of various local error estimators and indicators for adaptive grid refinement (indicators based

on the residual, second derivatives, Zienkewicz-Zhu, and hierarchical finite volume element methods). In this space, the research will work to develop a theory for error estimates of finite volumes of higher order, to study the corresponding strengthen Cauchy inequality and to develop the adaptive strategy.

### *Stabilization of the finite element method via least squares.*

The research will focus on the least-squares method for stabilizing finite element approximations of nonsymmetrical and indefinite problems. The main goal will be to address various stabilization approaches in solving convection-dominated diffusion equations. In collaboration with Panayot Vassilevski of CASC, the researcher will investigate stream-line diffusion least squares; the main tool will be multi-level preconditioning/realization of discrete inner product in Sobolev spaces with negative indices, multigrid methods and domain decomposition. Investigators also study elastic deformations of nearly incompressible materials in 3-D, and adaptivity in the context of the least-squares. Additionally, the researcher will interact with CASC to explore the FOSLS philosophy for design and implementation of robust and reliable algorithms for some limiting cases of Boltzman equation.





## Subcontracts, ASCI Level 3

---







# Unconditionally Stable Lattice Boltzmann Algorithms for Hydrodynamic Simulation

*Bruce Boghosian*

*Boston University*

## Abstract

The term “lattice Boltzmann algorithms” for computational fluid dynamics describes the time evolution of a discrete-velocity distribution function at each lattice point. Moments of these distributions then yield the desired hydrodynamic variables. Conservation of mass, momentum and energy (if appropriate) impose various equality constraints on these distributions, and positively impose various inequality constraints. We propose to use methods of linear algebra to coordinatize the subspace of distribution values that obey the conservation laws, and to use the Fourier-Motzkin elimination method to sequentially bind the coordinates within the polytope thus constructed. By then constructing an entropy function that reaches a maximum within this polytope, and goes to  $-\infty$  on its boundaries, and by demanding that the collision process increase this entropy, we can construct unconditionally stable lattice Boltzmann algorithms. Limitations on Reynolds number are then due to a minimum value of the attainable viscosity. We propose that the applicability of algorithms constructed in this way for applications in computational fluid dynamics be investigated.

# Research on Accelerated Strategic Computing Initiative (ASCI) Applications

*David Keyes*

*Old Dominion University*

## Abstract

The objective of this work is to introduce algorithmic ideas and research software implementations (in the form of the PETSc and PME libraries) useful to the parallelization of structured and unstructured grid PDE applications, implicit and explicit. The goal is to assist with the automation and optimization of parallel ports of "clean" legacy applications.

The work will research the following algorithms in the context of Accelerated Strategic Computing Initiative (ASCI) applications

- Domain decomposition.
- Newton-Krylov-Schwarz (NKS) methods for nonlinear equations.
- Nonlinear multilevel methods.

The researcher shall know how to restructure existing data structures and how to extract the specific expertise and investment in a code, while leaving behind solver-related code that is obsolete in the parallel environment.

# Algebraic Multigrid (AMG) Software


*John Ruge*

*Front Range Scientific Computation, Inc.*

## Abstract

The purpose of this project is the development and extensive testing of robust Algebraic Multigrid (AMG) software for solving unstructured-mesh problems that are designed to run in both serial and parallel environments. The emphasis will be on testing existing AMG codes on the problems of interest to Lawrence Livermore National Laboratory (LLNL), especially time-dependent 3D elasticity/plasticity problems and 3D heat conduction–diffusion equations. Another major focus is the design of highly parallel AMG strategies. Several technical approaches to improving AMG robustness and parallelism have been identified and studied in the framework of problems of interest to LLNL. Two major technical achievements were the conceptualization of schemes.





## University Collaborative Research Program Proposals

---





# Software Infrastructure for Multi-Tier Implementation of Structured Adaptive Mesh Hierarchies

*Scott Baden*

*University of California, San Diego*

## Abstract

The Principal Investigator (PI) proposes to investigate and develop a scalable portable programming methodology and software tool for efficiently implementing structured adaptive mesh refinement on multi-tier computers, which are multi-computers comprising symmetric multiprocessor (SMP) nodes. The program will leverage the investigator's previous work with the KeLP system. This C++ class library provides hierarchical SMPD control flow to manage two levels of parallelism and locality.<sup>1</sup> The primary research issue addressed by the investigation is to develop a "communication aware" multi-tier load-balancing strategy that includes communication as part of the computational workload, and permits the application programmer to mask the latency and overhead of communication.

This research builds on the PI's previous experience with structured adaptive elliptic solvers running on MPPs.<sup>2</sup> The investigation will deliver a computational testbed permitting Lawrence Livermore National Laboratory scientists to explore portable, scalable implementations of adaptive mesh applications running on a variety of platforms of interest to the Laboratory. This investigation is timely, as it will provide a practical technique for effectively utilizing not only cluster of SMPs, but also high-end ASCI platforms such as the IBM ASCI Blue-Pacific machine. Current approaches based on employing MPI alone are not scalable. The PI will

collaborate with Jonathan May (CASC) to carry out performance analysis of the software infrastructure and to establish contact with potential laboratory users.

<sup>1</sup> KeLP is part of the thesis research of Ph.D. student Stephen J. Fink, who is funded by a DOE Computational Graduate Student Fellowship.

<sup>2</sup> This pre-existing collaboration involves Scott Kohn (CASC), Beth Ong (CASC), and John Weare (UCSD, Chemistry).

# Computational Fluid Dynamic Studies of Arterial Flow Disturbance Induced by Intravascular Stents

*Abdul Barakat*

*University of California, Davis*

## Abstract

Atherosclerosis is an arterial disease whose pathological complications, namely heart disease and stroke, are the leading cause of mortality in the industrialized world. In its advanced form, atherosclerosis leads to plaques that protrude into arterial lumens and form stenosis or even complete vessel occlusions, which obstruct blood flow and give rise to the pathological events. One common interventional procedure involves the placement of an intravascular stent, an expandable wire mesh structure, that is introduced into the diseased artery in a compressed state and is inflated at the stenosis or occlusion site to both restore blood flow and provide structural stability to the arterial wall. The major limitation to the success of this procedure, however, is restenosis, a complex and incompletely understood process by which plaques re-protrude into the vessel lumen within a period of a few months.

The placement of a stent in an artery mechanically damages the endothelium, the monolayer of cells lining the inner surface of all blood vessels. In vitro data indicate that the rate of endothelial repair after injury may be significantly slower in regions in which endothelial cells are exposed to relatively large fluid mechanical shear stress gradients, as occurs at the end points of flow separation zones. Therefore, flow separation in the vicinity of a stent may contribute to restenosis. The hypothesis driving our research is that the

occurrence of flow separation depends on appropriate hemodynamic matching between the stent design and the flow and geometric properties of the arterial segment in which the stent is positioned. We propose to test this hypothesis by studying the impact of various geometric and flow parameters on the occurrence of near-stent flow separation using computational fluid dynamic techniques. We shall use the spectral-element fluid dynamic code NEKTON to accomplish this goal. Our specific aims are

- To perform three-dimensional steady-flow simulations of the flow simulations of the flow field in the vicinity of an intravascular stent positioned within a straight rigid-wall arterial segment.
- To extend the simulations to physiological pulsatile flow.
- To perform three-dimensional steady and pulsatile flow simulations of the flow field in the vicinity of an intravascular stent positioned within a curved arterial segment.
- To incorporate certain aspects of arterial wall motion into the simulations.

The proposed collaboration with LLNL is critical to the successful completion of the research for two principal reasons

- LLNL's supercomputer capabilities will be needed for the solution of the equations in the some of the more computationally intensive simulations.
- Hydrodynamic simulation capabilities currently under development at LLNL may, in the future, allow the incorporation of additional physiological considerations into the aortic model including fluid-wall coupling. Furthermore, results of the proposed research will define possible future collaborations with the Center for Health Care Technologies at Lawrence Livermore National Laboratory. Of particular interest are the areas of actuation mechanisms for the deployment of miniature devices within blood vessels, and the application of imaging modalities to better characterize the effect of stent placement on the arterial wall.



# Parallel Multigrid Methods on Unstructured Grids for Scientific Computing

*Tony Chan*

*University of California, Los Angeles*

## Abstract

We propose to investigate efficient multigrid methods for elliptic problems on unstructured grids that are suitable for distributed and shared memory parallel computing architectures. Two aspects of multigrid will be studied in this work.

The first is the design of optimized multigrid libraries for elliptic problems that are robust with respect to the mesh size and the coefficient of the PDEs. The second is the performance of various multigrid algorithms on parallel computers. Particular emphasis will be placed on multigrid algorithms appropriate for solving the discretization matrices arising from various large-scale scientific computing problems such as plasma physics simulations, turbulence flow, ground water flow, and the like.

# Renewal of a Collaborative UCLA–LLNL Program on Parallel PIC Modeling of Semiclassical Quantum Models

*Viktor Decyk*

*and*

*John Dawson*

*University of California, Los Angeles*

## Abstract

We propose to continue a collaborative program between University of California Los Angeles and Lawrence Livermore National Laboratory on parallel computing applied to particle-in-cell (PIC) codes. We have been successful in modeling many particle quantum systems by combining a semiclassical approximation of Feynman path integrals with parallel computing techniques previously developed at UCLA for simulating plasmas. With continued support, we should be able to construct a complete multiparticle quantum mechanical modeler.

# Massively Parallel Computations Applied to Finite Element Analysis

*Tony Keaveny*

*University of California, Berkeley*

## Abstract

Detailed stress analysis of trabecular bone, which has a highly complex and heterogeneous microstructure and unknown tissue properties, will improve understanding of its mechanical behavior, and accelerate progress in a number of biological and medical studies. Our goal in this continuing study is to utilize a unique finite element analysis method to study the failure behavior of trabecular tissue. Typically, these finite element models have 1–100 million elements and cannot be analyzed in an efficient manner using conventional finite element techniques.

During the first two years of this grant, a custom finite element code was developed which allows large models with more than one million elements to be efficiently solved using parallel computational techniques. Our code was used to study the convergence behavior of the models, and later to predict the elastic properties of trabecular tissue. The current proposal is to continue refining this analysis method and to then use it to study the failure behavior of trabecular tissue. As with our previous work on elastic behavior of trabecular bone, using this analysis technique, our proposed studies will combine extensive mechanical testing data available from ongoing experiments at UC Berkley with results from the finite element analyses to ensure model

validation. This work will be beneficial to Lawrence Livermore National Laboratory, as it will establish computational techniques that can be used for a variety of future research in computational biomechanics. These research areas include the effects of aging, bone injury, and osteoporosis.

# Development of a Three-Dimensional Relativistic Particle-In-Cell Code for Studying the Production of Useful Electron Bunches Using Ultra-Intense Laser Pulses

*Warren Mori*

*University of California, Los Angeles*

## Abstract

This is a continuation proposal for the “development of a parallelized three-dimensional (3D) relativistic particle-in-cell (PIC) code for studying the production of useful electron bunches using ultra-intense laser-pulses.”

During the past year, we developed an object-oriented skeleton code for 3D parallel PIC algorithms, studied laser-wakefield excitation in 3D using a code with the local charge conserving algorithm of Marder (Langdon’s improvement), and studied the electron beam characteristics for laser-driven wavebreaking and optical injection techniques in 2D. We propose to incorporate the 2D and 3D algorithms into the object oriented skeleton code, to benchmark these and other algorithms, and to investigate the production of useful electron bunches in 3D and 2D.

# A Computational Investigation of Finite-Strain Plasticity Model for Crystalline Solids

*Panayiotis Papadopoulos*

*University of California, Berkeley*

## Abstract

The proposed investigation is a continuation of current research on the development of a Lagrangian finite element-based computational framework for the analysis of initial/boundary-value problems of macroscopic rate-independent elastoplasticity, based on the theory of Green and Naghdi. The proposed work will concentrate on incorporating microstructural effects into the computational formulation, with particular emphasis on single crystals and polycrystalline aggregates. Detailed constitutive models within the above theoretical framework will be formulated and numerically analyzed in collaboration with ongoing research at Lawrence Livermore National Laboratory. This research focuses on experimental measurements in finitely deforming polycrystalline solids.

# Method of Simulation for Localized Multiply-Scaled Condensed Materials

John Weare

University of California, San Diego

## Abstract

Many materials of technological importance are of mixed composition, lack symmetry, and display complex chemical bonding patterns. All of these factors complicate the prediction of their structural properties. To make matters worse, the simulation of physical properties of such materials often requires dynamic information (e.g., prediction of reactivity or high-temperature pressure behavior). Typically, bonding interactions between particles for these materials are strong and many-body. This means that equations of motion must be solved in detail to predict dynamics. Since atomic particles are massive, the dynamics can be simulated directly using Newton's laws, as in molecular dynamics (MD). However, materials with covalent, metallic, and even hydrogen bonding are difficult to describe succinctly with simple analytical inter-particle potentials, making them difficult to simulate with traditional MD approaches.

Recently, considerable progress has been made with the introduction of *ab-initio* molecular dynamics (AIMD) methods. These method are similar in intent to traditional MD simulation methods, except that the forces necessary to propagate the system are calculated from very fast solutions of the electronic Schrodinger equation within the Local Density Approximation (LDA). This greatly generalizes the application of MD and has made possible considerable progress in material simulation. However, there are still very significant limitations to the application of these approaches, particularly for systems

with highly bound electrons such as O, F, and the transition metals.

The research program proposed here is designed to both develop new methods of simulations targeted at these difficult elements and make progress in the actual simulation of materials of direct interest to research programs at Lawrence Livermore National Laboratory (LLNL). The progress in simulation will come from the more efficient parallel implementation of existing methods and from the development of new simulation approaches. The proposed methods will more efficiently resolve the electronic wavefunctions in regions of rapid variation and, thereby, focus computer memory and CPU cycles in the portions of the computational domain where they are most needed.

We will pursue two methods. The first represents a continuation of our previous efforts to implement a structured adaptive finite element method that places structured grids of increasing refinement in regions requiring higher resolution (e.g., near atom centers). We have already made some progress with this approach; however, we have found that we need to go to more accurate discretizations to obtain the required chemical accuracy. We are hoping that chemically accurate results can be obtained by a local application of spectral element methods. The second approach to local enhancement is to use a mixed basis set of planewaves augmented by a linear local basis confined to regions near the atomic centers. This would amount to a dynamic implementation of LAPW methods. LAPW methods are known to perform well for these systems. However, it will be a challenge to develop a parallel algorithm efficient enough to allow dynamics for sufficiently long simulation periods.


One objective of LLNL's H Division is the development of simulation methods capable of predicting high-temperature high-pressure behavior in the system  $HF - CO_2 - CH_4 - H_2O - N_2$ . For the temperatures and pressures of interest ( $T \approx 3400$  K and  $P \approx 32$  GPa), there is very little data available, even for the pure end members. This is particularly true of the HF system because of its corrosive nature. The intention of this program is to replace unavailable experimental data by computational data for moderately sized systems. The presence of O, N, and F make this a difficult problem for AIMD. However, preliminary calculations have shown that HF can be

calculated with high accuracy by choosing a softened pseudopotential. Since the number of valence electrons in each of these molecules is relatively small, we expect that we will be able to simulate systems of the order of 40 molecules with our present algorithm. This will require fast implementations of our codes on LLNL's parallel supercomputers, and the efficient utilization of memory hierarchy. We believe that the present solver algorithms can also be improved by preconditioning the conjugate gradient iterative method. Furthermore, because the intermolecular interactions in this system are relatively weak (hydrogen bonds), we will need to implement and test density correction approximations. All of this can be accomplished in the first year of the grant.

As we make progress with the new methods, we will use  $HF - CO_2 - CH_4 - H_2O - N_2$  as a test of efficiency and accuracy. The proposed computation approaches may be implemented so as to eliminate the requirement of pseudopotentials. All-electron calculations of this sort may be required to treat very high-pressure systems. We emphasize that while the methods we are developing as part of this grant will be used to treat problems in the  $HF - CO_2 - CH_4 - H_2O - N_2$  systems, these new methods will also have general application to other problems that involve difficult elements, such as transition elements.







## LDRD Projects

---





# A New Approach to Orthopaedic Implant Design

*A Laboratory Directed Research and Development Proposal for Exploratory Research in the Institutes*

*Karin Hollerbach*

## Abstract

Our approach to implant design, which uses finite element modeling (FEM) of joints with implants, offers an opportunity to evaluate designs before they are manufactured or surgically implanted. This is an attractive alternative to current methods that use only experimental testing for design evaluation. Failures depend on implant and human tissue geometry, the constitutive properties of the biological and the implant materials, and/or interaction between the human tissues and the implant components. In all cases, FEM can contribute to early changes in the design process, resulting in better, less costly and longer-lasting implants.

# Novel Parallel Numerical Methods for Radiation and Neutron Transport

*A Laboratory Directed Research and Development Proposal for Joint Exploratory Research in the Institutes and Exploratory Research in the Directorates*

*Peter Brown*

## Abstract

The availability of advanced computational methods for transport modeling is of crucial importance to the Department of Energy and to Lawrence Livermore National Laboratory. In many of the multi-physics simulations performed at the Laboratory, transport calculations can comprise from 30-50% of the total run time. Thus, a significant core competence in the formulation, software implementation and solution of the numerical problems arising in transport modeling is an essential component of the Laboratory's computational repertoire.

The goal of the proposed research effort is to enhance these capabilities through the design of advanced numerical methods for the parallel solution of 3D radiation and neutron transport problems. The equations describing the above processes are usually some variant of the Boltzmann transport equation (BTE), which requires the solution of very large systems of equations. The ability to solve such problems efficiently requires a combination of several kinds of numerical methods for parallel computers.

Recent developments in the area of first order system least squares (FOSLS) methods show great potential for

providing more accurate and robust solution procedures than current approaches. The FOSLS-based approach also provides a natural multigrid solution procedure for the resulting discretized problems. We also propose to develop more accurate phase space discretization techniques, in space and direction (angle) as well as better time-stepping algorithms is also of interest due to the absence of ray effects in spherical harmonics solutions. The harmonic projection algorithm we have developed in earlier work uses discrete-ordinates codes to obtain spherical harmonics solution, and we will investigate its usefulness in solving Laboratory problems.



## Students and Faculty

---





# Developing New Scalable Solver Methods

*Mark Adams*

*University of California, Berkeley*

## Abstract

Mark Adams continued to develop scalable solver methods for large-scale unstructured finite element problems in solid mechanics. He originally implemented these methods in the course of his research towards his doctoral dissertation at UC Berkeley. During the summer of 1998, he worked with the Linear Solvers Group in the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL). His research at CASC focused on scalable solver efforts with regard to the software design of parallel linear solver libraries, and construction of scalable numerical primitives, common to his solver and that of the algebraic multigrid group in CASC. Mr. Adams also continued to discuss issues related to the design and use of PETSc (from Argonne National Laboratory) for the multigrid codes in CASC.

He continued to develop solution algorithms for contact problems with Lagrange multipliers, via Uzawa type algorithms, for the ALE3D group at LLNL. Mr. Adams has a serial prototype within his parallel code (using perfectly scalable constructs) that effectively solves some of the simple contact problems from the ALE3D group.

At the conclusion of the summer, Mr. Adams was able to solve problems exhibiting up to 9.6 million degrees of freedom in large-deformation, incompressible elasticity (Poisson's ratio of 0.49) and plasticity problems, with large jumps in material coefficients ( $1.e-4$ ), on a Cray T3E and an IBM PowerPC cluster.

# Change Management in the Digital Library Environment

*Igg Adiwijaya*

*Rutgers University, Newark, New Jersey*

## Abstract

Igg Adiwijaya worked with the DataFoundry project, specifically in the area of change management of autonomous data sources. DataFoundry is a data warehousing and integration research effort underway in the Lawrence Livermore National Laboratory Center for Applied Scientific Computing (CASC). Mr. Adiwijaya familiarized himself with DataFoundry and the sources from which the project extracts data. He also researched the methods for detecting and extracting changes to external data sources. He then evaluated the available approaches, including their advantages, limitations and potential for applicability to DataFoundry. He also implemented an automated process for extracting data from external data sources, and scheduling periodic runs.

Mr. Adiwijaya proposed a general framework for change management in the Digital Library environment, particularly the one serving DataFoundry. For DataFoundry, the framework requires approaches such as comparing two semi-structured documents, ranking semi-structured documents based on schema, and extracting changes to semi-structured documents. He devised an effective and efficient approach for meeting such requirements.

While concluding his summer with ISCR, Mr. Adiwijaya conducted further research on representation of changes and data mining on the extracted changes. He also started the implementation of the parser and delta extractor, components in the framework, and prepared one technical manuscript documenting the change management framework.



# First Order System, Least Squares (FOSLS) Methods for Radiation and Neutron Transport

*Travis Austin*

*University of Colorado, Boulder*

## Abstract

Over the summer of 1998, we began to implement an idea for accurately describing neutron transport within different types of media. Many current ideas are based on methods that suffer in particular types of media. For example, researchers currently are forced to deal independently with problems defined in the types of media exhibiting the property that neutrons have a small mean free path before collisions occur. We desire to treat such problems with the same ideas we are using in other cases. More importantly, we want our method to be successful, independent of problem characteristics.

In designing the problem formulation, we insisted that our solver have certain characteristics. This drove us to pose the problem as a least squares minimization problem. We coupled this idea with an expansion of our solution in spherical harmonics and a discretization of the moments with finite elements. To accurately treat all cases, we chose a scaling of the matrix equation based on the problem characteristics. Since we want a robust solver, we considered using multilevel techniques that have the potential to achieve convergence factors that are independent of mesh size.

The backbone of the described transport code has been established. Yet, we discover that our desire to use a multilevel solver and be robust compels us to pay special attention to solving for the first order moments. Multilevel solvers are extremely efficient and accurate when applied to certain classes of problem. Unfortunately, our

formulation creates a system of equations for the first order moments that standard multilevel solvers do not handle well. This requires treating this case individually.

The second half of my summer dealt exclusively with this addressing this problem. We wanted to show that we could get good convergence results for the first order moments. With this established, we would be able to move on confidently to establishing the success of our method for the more general neutron transport problem. These characteristics were evident in our code, and this demonstrated that handling the first order moments as independent of the other moments was necessary. The code designed also showed that we could be successful in this endeavor.

# Examining the Literature on Colloid and Colloid-Facilitated Transport Modeling

*Lora Ballinger*

*University of Utah*

## Abstract

Current literature on colloid transport and colloid-facilitated transport was examined to evaluate the state of knowledge and the strengths and weaknesses of available models. Among the journals included in this literature examination were *Colloids and Surfaces*, the *Journal of Contaminant Hydrology*, the *Journal of Nuclear Science and Technology*, *Nuclear Technology*, *Radiochimica Acta*, *Radiochemistry*, and *Water Resources Research*. This literature search also included articles from conference proceedings (Scientific Basis for Nuclear Waste Management Symposia, Materials Research Society Symposia), and technical reports from LLNL and the Center for Nuclear Waste Regulatory Analysis. An annotated bibliography of 51 entries, which analyzes the relevance of the literature examined, was produced.

In addition, the relevant points from the literature were detailed in a report. In summary, two types of colloid must be considered: pseudocolloids, radionuclides bound onto natural colloids; and radiocolloids, composed entirely of radionuclides. Models are needed to describe the formation of each of these types. The typical colloidal transport model requires knowledge of parameter values that are not easily measured. In addition, in some cases, the theoretical predictions are insufficient to describe observed behavior; therefore, models based on this theory will not be accurate. A colloid velocity model must be developed that takes into account hydrodynamic chromatography effects. In addition, filtration, particle size and particle chemistry confer major effects on the distance traveled and the velocity of colloidal particles; these effects should be explicitly included in the model formulation. Finally, very few of the models in the literature have been validated adequately, particularly in field situations.

Based on information in the literature, recommendations were made regarding the necessity and feasibility of incorporating a colloid-facilitated transport model into the reactive transport model of the Yucca Mountain system. In some cases, the literature indicates that colloids could possibly have a significant effect at the Yucca Mountain site and should be explicitly included in the model. Most studies, however, consider colloid transport in a porous medium and do not consider the case of episodic flow in a fracture. Therefore, one cannot predict with any accuracy the effect of the colloids without including them in the model. This problem is further compounded by the fact that a majority of the data needed to predict parameter values for existing models is not available — either at all or for this site specifically.

It was recommended that an effort be made to obtain information on colloid diffusion coefficients and filtration constants (may involve surface potentials of the fractures, as well as density of binding sites and electrolyte concentration). Recommendations were also made for specific components of the model: As few colloid formation models exist, it was suggested that a simple equilibrium model could best serve for the time being. Particular functional forms for adsorption onto and desorption from colloidal particles were given. A simple means for combining two of the typical models for colloid velocity was presented. A basic form of the general transport equation was given, and a list of significant factors to was given. Among these significant factors are the distribution of fracture apertures and particle sizes, and the density of binding sites on a particular material.

Topics of interest in developing a model, including pseudo-colloid formation, filtration effects, and particle size and chemistry effects, were addressed. The availability of model parameter values at the Yucca Mountain site was determined by a literature search. Finally, a simple model of colloid transport was developed and implemented using the available parameter values. The model used a finite differences formulation to describe movement (not formation) of colloids within a 1D column. Results from the model runs indicate that, using the values gleaned from the literature, it would take only a couple of years for a particular location on the fracture to saturate, once the colloids arrived there. In addition, it was seen that by 10 years, the walls of a fracture 40 meters long would be completely covered. These could be good indications, or bad. If the number of binding sites were unlimited, quick binding would be good; radionuclides would be trapped on the walls of the fracture. If the number of binding sites were small, and the colloids could bind to each other, as well as to the fracture walls, then they could plug the fracture and prevent further transport. The results otherwise indicated, however, that the particles move very quickly in comparison to dissolved particles.

# Research Progress in Pattern Classification, and Seminars Given on Object Recognition, Neural Networks, and Genetic Algorithms

*George Bebis*

*University of Nevada, Reno*

## Visiting Faculty Abstract

During my visit to LLNL and CASC (July 8, 1998 through August 20, 1998), my major responsibility was to assist Sapphire Project members in identifying the preprocessing and pattern classification techniques appropriate for the FIRST image data (from the survey of Faint Images of the Radio Sky at Twenty centimeters). We concentrated on the problem of identifying radio sources exhibiting the special morphology known as "bent-double." This problem is quite challenging, as astronomers have a hard time defining bent-double morphology precisely. They also experience difficulty in confidently classifying a radio source image as a bent-double. I believe that many researchers find it very hard to come up with a general and complete definition for characterizing bent-double morphologies. This problem is thus very suitable for pattern classification techniques that extract knowledge about the problem at hand by examining a large number of examples (e.g., neural networks, decision trees, and the like). For these techniques to be successful, however, it is vital that we provide researchers with important information (features) extracted from the examples. This requires the use of sophisticated noise reduction and segmentation techniques.

In addition to my Sapphire Project responsibilities, I also gave one formal and two informal seminars. In the first, I reviewed the problem of object recognition, and then presented my recent work on the application of algebraic

functions of views for object recognition. In the second seminar, I presented the fundamentals of neural networks, and then discussed an interesting application involving the use of neural networks for face detection from gray-scale images. In my final talk, I presented the fundamentals of genetic algorithms and discussed my recent work on the use of genetic algorithms for object recognition.

## Work Completed

During the first phase of my work, I focused on the problem of feature extraction. My goal was to develop approaches for extracting features in a way that will allow pattern classification techniques to separate bent-double from non-bent-double morphologies. My activities are summarized in the following bullet points.

- **Appropriate image processing, neural networks, and machine learning software.** I did an extensive search to identify appropriate software to be used in the project. I also helped install and test some of the software packages.
- **Astronomical Image Processing System (AIPS).** AIPS is a popular software package for processing astronomical images. As I had no familiarity with this package, I spent some time trying to understand which of its functions might be useful in extracting important features from the radio source images. I also spent some time trying to understand FITS, the format in which astronomical images are stored.
- **Information stored in the catalog of the FIRST survey.** Astronomers who performed the FIRST survey have already extracted and catalogued some of these images. I spent a good amount of time trying to understand what information had been extracted from the images and how it can be useful for pattern classification. My conclusion is that the information stored in the catalog will be insufficient in helping to classify a radio source as a bent-double or alternative morphology. Additional features should be extracted (see Feature Extraction below). The information stored in the catalog, however, will help us identify promising regions of the sky, that is, regions where it is very likely that radio sources can be found. Concentrating on these regions reduces the time needed to search the whole sky.
- **Reviewing the literature.** I spent a good amount of my time reviewing the literature. I found and studied many interesting papers on the application of classical pattern classification, neural networks, and machine learning techniques in astronomy, as well as in other related areas. Neural networks and

machine learning techniques generally have been applied to the problems of galaxy and star classification. Good results have been reported in these cases, which lends me to believe that they should also work well for our problem. Copies of these papers were provided to other Sapphire investigators.

- **Noise reduction and segmentation.** Before features can be extracted from the images, it is important to filter out the noise, apply some kind of image enhancement, and then separate the regions of interest from the background. During the limited period of time I stayed at LLNL, I experimented with several noise reduction and image enhancement techniques, and I came up with a procedure that seems promising. I implemented and successfully tested this approach. Chandrika Kamath and the other members of the group have a good idea about how this approach works.
- **Feature extraction.** After the regions of interest have been isolated, the next step is to compute features within those regions that might be useful for classification. Since it was not obvious to me which features were the most important for our problem, I decided to extract as many features as possible. In my current implementation, the features I extract are based on geometric properties of the regions (area, compactness) as well as histogram features (texture, moments).
- **Feature selection.** Presenting all possible features to the pattern classifiers would very likely overwhelm them, particularly as some of these features carry no discrimination power benefiting the problem we are trying to solve. This technique would yield very poor results. Thus, it is essential to first apply feature selection—techniques that reduce the dimensionality of features before they are fed to the classifier. A popular method known as principal components analysis (PCA) considers linear combinations of features. I implemented PCA, but did not have time to test it during my stay at LLNL.
- **Classification.** I designed a method for distinguishing bent-double from non-bent-double morphologies based on neural networks. Due to my limited visit, however, I did not have the time to implement and test it. I discussed this method with the other members of the Sapphire Project, as well as with Bob Becker from the LLNL Institute of Geophysics &

Planetary Physics (IGPP). He finds it interesting and believes it is worth trying.

### Work to be completed

- **Parameters for noise reduction and segmentation.** The procedure I have implemented for noise reduction and segmentation can give very good results but it depends on some parameters. If these parameters are not chosen appropriately, the segmentation results can be very poor. I implemented a number of heuristics to choose these parameters adaptively; however, more heuristics and much more testing are needed in order to come up with a scheme delivering very good performance. I should emphasize that the noise reduction and segmentation steps are very important and will influence the performance of the whole system if the results they produce are not very good.
- **Feature extraction, selection, and normalization.** We should consider extracting additional features that can provide more information to the classifiers. Additionally, the PCA approach as well as other approaches must be applied for selecting only the most important of the features we intend to feed to the classifiers. Issues related to the normalization of the features must also be considered, especially when neural networks are used for classification.
- **Non-bent-double training examples.** Bob Becker and his group have already provided us with images of radio sources that likely correspond to bent-doubles. However, in order to train a classifier to distinguish bent-doubles from non-bent-doubles, a representative number of non-bent-double examples must be included in the training set. Careful selection of the non-bent-double examples is critical in minimizing false positives as well as false negatives.
- **Appropriate classification schemes.** A large variety of classification techniques is available for investigators to try on a given problem. I believe that after we have come up with a good set of features, it becomes important to test various neural network and machine learning techniques. Based on my experience, finding the best technique for a problem as well as fine-tuning it (e.g., choose the appropriate number of hidden nodes in a neural network) requires a lot of experimentation. Recently, many researchers are interested in combining the results of several different classification techniques (fusion) in order to improve performance. Results based on this idea have been promising, and I believe that we should also consider this direction.

# Wavelet-Based Compression Schemes

*Martin Bertram*

*University of California, Davis*

## Abstract

I have implemented a wavelet-based compression scheme for time-varying volumetric data sets. For the wavelet transform, wavelet lifting, in-place computation of the coefficients and integer arithmetic were applied. After several experiments it turned out that a linear B-Spline wavelet performed best in decorrelating the data for compression.

I also developed an efficient coding scheme. This scheme can be further improved by taking into account the hierarchical structure of the wavelet transform and different probability models for each single coefficient or each bit of a coefficient.

# High-Performance Surface Rendering

*Kathleen S. Bonnell*

*California State University, Dominguez Hills*

## Abstract

I worked with Dan Schikore on high-performance surface rendering. In particular, I investigated performance differences between different triangular models. My first tests involved the SGI Infinite Reality Engine benchmark program, which renders 10 million triangles per second using a triangle strip representation of the data set. Currently-render performance of individual triangle models yields only about 10% of that stated performance of the architecture. For my tests, I measured how varying the length of the triangle strips affected render performance. The motivation behind these tests was to increase performance. Individual triangles have a cost of 3 vertices per triangle, whereas triangle strips have a number of triangles plus 2. It would seem to follow that the longer the strip one could create, the better the performance should be.

My tests indicated otherwise. In fact, performance peaked at a range of 10–12 triangles per strip; after that it leveled off, and longer strips did not help to improve performance. On the flip side, the test results dramatically illustrated the fact that strips of length less than 10 were a definite hindrance to performance.

After establishing this, I moved on to tests involving generalized triangle strips (GTS). Conceptually, generalized triangle strips can turn corners. Thus, one can generate extremely long strips by turning enough corners. Old GL hardware supported GTS via a swap facility with minimal cost. Newer GL hardware can only simulate a swap by sending a duplicate vertex to the renderer, with an obvious cost of one vertex per swap. We wanted to test the performance of strips with swaps as opposed to strips without swaps on the same data set. To that end, I modified an existing software package

(Stripe by Francine Evans, State University of New York at Stony Brook). I modified the package to work with isosurfaces and material boundaries for ASCII data. I also created a no-swaps version of the software. Results of performance tests with this modified software package demonstrate that swaps performed better than no-swaps; and minimizing swaps performed better than not minimizing.

At first these results might be surprising. The Stripe package, however, was designed to create the longest strips possible, regardless of the number of swaps involved. When I removed the swaps from the strips, this created a great number of strips of length less than 10. Based on the tests with the SGI benchmark, the poor performance of my no-swaps version is understandable.

The next step would be to test an algorithm that creates strips with a length of at least 10 triangles— both a swaps and a no-swaps version. Unfortunately, I ran out of time, so that must be left for a future project.



# Algebraic Multigrid Visualization

*Tim Chartier*

*University of Colorado, Boulder*

## Abstract

For twelve weeks during the summer of 1998, I worked as a summer student at Lawrence Livermore National Lab (LLNL). In particular, I worked within the Center for Applied Scientific Computing (CASC). Van Henson and Jim Jones guided my research.

A focus of the multigrid research team at LLNL is to apply algebraic multigrid (AMG) technology to a host of problems of importance to the Lab. CASC intends to further develop AMG and to find parallel algorithms.

My summer began with work on AMG visualization. To aid in research, it is often helpful to see the grids that AMG produces and to plot the residual for the cycles. Using Matlab, I produced pop-up windows to visualize both of these important aspects of AMG. Others and myself used this tool throughout the summer. The tool was designed for my work in AMGe, which is a new version AMG based on element interpolation. This tool can easily be transported into existing AMG code used at LLNL.

A goal of the multigrid research team at LLNL is to numerically solve elasticity problems. AMGe shows promising results in such problems. Yet, an open question is how to find “good” grids for AMGe. In two-level multigrid, we want to solve for the fine grid error on a coarse grid, and then interpolate this error back to the fine grid. My goal was to begin research on such coarsening.

AMGe offers new tools for deciding what sets of points would give good interpolation. These tools supply large amounts of new information. How could this information

help us find numerically useful grids? With the help of Jim Jones and Van Henson, I implemented an algorithm to begin to answer this question.

As the summer progressed, we gained valuable insight on our new AMGe tools. Useful discussion brought new results. These results further developed options and ideas for future research. By the end of summer 1998, I left LLNL with many ideas for further research in coarse grid selection for AMGe. AMGe shows promise as a numerical solver for difficult and important problems such as elasticity. Such an algorithm would help LLNL with its projects and applications. AMGe exhibits promise both in its usefulness and its robustness.

# Rayleigh-Taylor Mixing: Computational Investigations of Linear Electric Motor Experimental Data

*Jonathan Dursi*

*University of Chicago ASCI Alliance<sup>1</sup>  
Tri-Lab Alliance Technical Sponsor: Stephen Libby*

## Abstract

At the heart of understanding the accuracy of interface mixing simulations lies the careful validation of numerical models with correctly interpreted experimental data. In this study, Jonathan Dursi performed a variety of geometrical and statistical studies of data from linear electric motor (LEM) fluid interface experiments by Dimonte et al. at Lawrence Livermore National Laboratory). Such studies are a necessary prerequisite to any serious experiment-theory comparison.

The LEM data consists of 2D slices through a 3D Rayleigh-Taylor mixing zone, with two immiscible fluids of Atwood number of 0.34 and subject to an acceleration of 75 G. In each experiment, images from three time frames show the deeply non-linear evolution of the mixing zone. To understand the mixing physics, the approach—rather than to examine various spectra of the resulting density field—was to investigate individual “blobs” for spikes of heavy fluid penetrating the lighter fluid, and “bubbles” for light fluid intrusions into the heavier fluid. The images from the experiments had already been processed (smoothed and bi-leveled) into values which represented either the high- or low-density fluid. The blobs could then be easily identified and tabulated by performing a linear search through the images, using a flood-fill algorithm to identify connected regions. This allowed some very simple investigations to be performed.

Dursi began by examining the evolution of the number of blobs as a function of the penetration depth. Although many mixing models typically use very peaked distributions of blobs, the distribution was clearly seen to quickly become quite flat. Similarly, he investigated the size distribution of the blobs, and the evolution of the distribution. The distribution again seemed remarkably flat, but with a strong peak for small blobs; however, “slicing” effects through the 3D experiment will tend to bias results toward small, observed blobs. It is not yet clear to what extent this bias affected the results.

More sophisticated analysis is underway, involving modeling the observed blobs with simple shapes, and noting the resulting size/shape distribution over time under constant and varying accelerations. This analysis, as above, requires untangling the effects of taking a 2D slice through 3D objects. Further work on this project involves comparing experiments with simulations, and examining more carefully the image processing procedure to determine sensitivity of results. When 3D simulation results are available, the experimental procedure for image acquisition and processing will be followed within the simulation cube, to both help understand the experimental data, and to potentially verify the simulation results. Sensitivity analysis will be done by examining sensitivity of results to varying the procedure and parameters of the image post-processing.

<sup>1</sup> The University of Chicago ASCI Alliance is focussed on advancing the understanding and modeling of astrophysical thermonuclear flashes that are due to the accretion of material onto a compact body and subsequent ignition. These flashes manifest themselves respectively as x-ray bursts, type 1a supernovae, or novae depending on whether the accretion is onto a neutron star or a white dwarf (with either an off-center carbon flash for a supernova 1a or a thin shell hydrogen flash for a nova). These flashes are significant in the production of light elements, in the accurate determination of the cosmic distance scales, and in probing the physics of the compact objects themselves. Since the accumulation of material and its ignition mode depends crucially on turbulent mixing, a significant part of the effort at Chicago is focussed in this area. In particular, the effort of the two LLNL summer students, Jonathan Dursi and Yuan-Nan Young, was aimed at advancing and validating Rayleigh-Taylor interface mixing models, layer formation, and subgrid turbulence models.



# Multigrid for Problems with Local Refinement

*Christopher Higginson*

*University of Colorado, Boulder  
Department of Applied Mathematics*

## Abstract

To obtain efficient multigrid solvers for problems with discontinuous coefficients, generally one must take the problem coefficients in to consideration in defining the components of the multigrid algorithm. For example, in the so-called “black box” multigrid method, the operators used to transfer functions between grids are defined based on the coefficients of the fine grid matrix. at the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory, these black-box operator-based interpolation formulas have been used successfully in two parallel multigrid solvers, SMG and ParFlow.

Applying these operator-based interpolation formulas when the global grid is logically rectangular is relatively straightforward. Together with Jim Jones of CASC, Chris spent the summer looking at methods for applying these black-box multigrid ideas to grids with local refinement, where the global grid will no longer be logically rectangular. In particular, he looked at refinement for cell-centered discretizations.

Chris looked in the literature for existing robust cell-based multigrid solvers. He implemented several of these and tested them on difficult diffusion problems with discontinuous coefficients. Based on their lack of robustness, Chris developed a new algorithm for cell-based multigrid. By careful construction of an interpolation operator that took into account variation in the diffusion coefficients, he was able to improve multigrid performance. The new algorithm was successful on several test problems where the existing methods had failed. Together with Jim Jones, Chris wrote a draft paper describing the new method. This work should provide

some valuable information about how to transfer residuals and function values at the fine-coarse interface, an issue that must be addressed as researchers build general purpose multigrid solvers for use in codes with structured adaptive mesh refinement.

Chris also spent part of the summer looking at a first order system, least squares (FOSLS) approach to parabolic problems. This approach has the potential to provide rigorous error estimates to guide refinement in both space and time. Chris developed the discretization method and began looking at parallel space-time multigrid solvers for the resulting discrete equations.

# A Problem-Solving Environment for the LOCFES Suite

*Rithea Hong*

*Texas A & M University*

## Abstract

My project for the summer, under the supervision of James McGraw, was to design and implement a problem-solving environment for the LOCFES suite. This suite is a group of one-dimensional neutron transport codes that my thesis advisor, Dr. Paul Nelson, has developed over several years. These codes solve the monoenergetic, steady state, one-dimensional neutron transport problem in plane parallel geometry. They are tools designed to aid in the study of various spatial discretization schemes. The main elements of my project were:

Identify essential elements of an environment that would make using these codes easier and more efficient.

Design the environment.

Implement the environment.

At the beginning of the summer, I had only a conceptual idea of what the environment should contain. My work began by identifying what would make the user's life easier and more efficient. The existing suite of codes operates by reading in an input file containing all of the relevant parameters, and then, eventually, producing an output file containing the results. Although it is simple to code, such an interface is a barrier to the user. The user must first learn the details of the syntax for the input file before using the code. Also, once the user has learned the code, he or she must handle the management of input and output files individually. All this is tedious to the user and wastes time. I designed an environment that would provide for interactive problem specification, result analysis, and management of both input parameters and results. During the design phase, I obtained feedback from several people, including those

familiar with neutron transport and those familiar with GUI design. Their input proved extremely useful.

Before implementation began, it was necessary to survey the available development environments and to weigh carefully each of their strengths and weaknesses in terms of what the project required. I felt that the most desirable development environment would be sufficiently powerful to implement the designed system, be sufficiently mature so that defects in the system would not interfere with the implementation, be accessible so I could actually use it, and be widely used. The final requirement ensures that an infrastructure exists to aid during development, and later maintenance. One does not want to choose an orphaned environment. After weighing several options, I concluded that Java was the best option. Since it was designed in the era of GUIs, it provides powerful support for the rapid development of the environment I envisioned.

After selecting Java, I spent some time familiarizing myself with the language before beginning implementation. After becoming comfortable with Java, I began implementing the GUI portion of the environment. By the end of the summer, the GUI portion was largely complete. Remaining tasks include: implementing the storage and retrieval of parameters and results, integration with the LOCFES code, and analysis of the results.

During this work, I also attended numerous presentations and interacted with a wide variety of Livermore personnel. I thoroughly enjoyed the experience and made considerable progress in my project. Particular thanks go to James McGraw, Neale Smith, Neil O'Neill, and Milo Dorr who helped me with my projected throughout the summer.

# Amino Acid Descriptor Clustering Tools

*Rachel Karchin*

*Stanford University*

## Abstract

During the summer of 1998, Rachel Karchin participated in the Lawrence Livermore National Laboratory summer student program by working with researchers in the Center for Applied Scientific Computing (CASC). She wrote Perl scripts to implement a web page that queries the DataFoundry warehouse. The web page visitor selects from a wide variety of attributes and conditions available on the page, and the scripts produce SQL query statements, execute them, retrieve data from the server, and display the results in tabular format.

Karchin also created several new html pages and Perl scripts to implement amino acid descriptor clustering, using the DataFoundry warehouse. These tools enable a user to:

- View all the amino acid descriptors in selected pdb chains.
- Perform computations on descriptor columns.
- Create new clusters of amino acid descriptors by selecting and launching a clustering program.
- View clustering results and write them back to the warehouse with additional annotation, the date of the experiment and a description, entered by the user.
- View the most probable sequence within a user-selectable window size for an existing cluster of descriptors.

Additionally, Karchin researched available machine learning public domain software for the Sapphire project and installed and documented the decision tree programs oc1, ripper, and cn2 and the clustering program autoclass. Independently, she studied C++ and machine learning; took several Lab TV courses, including Java Projects, C++ iostream, and Calculus and Probability

Refresher for Engineers; and attended object-oriented design users group meetings.

Before leaving CASC, Karchin gave a 45-minute presentation on hidden Markov models and their applications to protein classification and multiple alignment. She presented this material to a group of Lab Biology and Biotechnology Research Program (BBRP) scientists.

Karchin demonstrated the SAM software suite to a group from BBRP and CASC's DataFoundry Project. SAM was written at UC Santa Cruz, where Karchin was one of the coders. It implements hidden Markov model technology for use with protein, DNA and RNA sequences. She also installed a local copy of SAM on one of BBRP's SGI servers.

# Visualization of Large Scientific Data Sets Using Virtual Reality

*Falko Kuester*

*Center for Image Processing and  
Integrated Computing (CIPIIC)  
University of California, Davis*

*email:kuester@cs.ucdavis.edu  
graphics.cs.ucdavis.edu/~fkuester*

## Abstract

This summer research project, conducted at Lawrence Livermore National Laboratory (LLNL), in the Center for Applied Scientific Computing, focused on two distinct topics. The first was real-time visualization of large scientific data sets, using optimally adapting meshes and virtual reality (VR) display technology. The second was the motivations and benefits for the use of VR as a possible next-generation visualization and modeling environment. The requirement and implementation ideas for adaptive meshes are noted below.

## **Interactive Visualization of Large Data Sets**

The key idea of this project was to identify an appropriate modeling scheme necessary for the visualization of huge complex scientific data. The Real-time Optimally Adaptive Meshes (ROAM) algorithm developed by Mark Duchaineau and others was selected as the underlying design idea and possible improvements were investigated. Our primary interest was to identify a real-time-capable surface or volume model satisfying the specific performance requirements listed below. We considered a variety of models, including parametric, implicit and polygonal representations.

## *Requirements*

The following performance and feature requirements were established

- General surfaces (e.g. ISO-surfaces).
- General viewpoint notion.
- Highly agile view direction.
- Parallel run-time optimization.
- Improved run-time efficiency.
- Incremental frustrum culling.
- Visibility culling.
- Time dependent surfaces.
- Topological changes vs. scale.
- Texture hierarchies for general surfaces.
- Performance comparison to estimated true optimum.
- Better performance model.
- Geometry paging from disk (preferably with compression).
- Texture paging from disk (preferably with compression).
- Non-pre-computed adaptation of meshes.
- Multi-pipe display.

Related issues included hardware model for performance optimization, modifiable display lists, and parameter controls.

## *Results*

The evaluation of surface models indicated that implicit surfaces and certain parametric representations like the Cao En surface model promised the best results. Demo applications have been created for these surface types.

## **Visualization using VR Technology**

Problems associated with the visualization of the massive scientific data sets generated at LLNL demand new technologies. As part of this project, a feasibility study of virtual reality as a next generation visualization environment was launched. Preliminary results are discussed at the end of this section.

## *Terminology*

In the early Seventies, Myron Krueger defined the term artificial reality, which was followed by derivatives such as cyberspace, virtual reality, virtual environments (VE) and synthetic environments (SE). The last term was

defined in 1992 by the Committee on Virtual Reality Research, established by the National Research Council. The terms VR and SE are used interchangeably in the following sections.

### *Vision*

Push your keyboard to the side, put on your data gloves, turn on your 3D display and engage a truly 3D visualization and modeling environment. Either visualize pre-computed results or start from scratch by modeling the simulation environment itself. Beginning with geometric primitives or a pile of “virtual clay,” and with the help of virtual design tools, create the initial model for a fluid mechanics or thermodynamics simulation. Communicate with other scientists through a centralized database-oriented design-space giving instant access to the latest 3D visualizations. Verify the setup and boundary conditions by running a few simulation cycles and visualize them in the created environment before spending hundreds of CPU hours on a problem with incorrectly defined boundary conditions. This mechanism would allow scientists to turn their insight into concepts, concepts into models, and finally models into production data or research results. The objective is to develop a suite of tools that will dramatically reduce cost, time, and complexity in research and development.

For the initial study, the focus was on the visualization aspects. The additional modeling component in this vision could be a powerful application area for the modeling schemes investigated.

### *Requirements*

Synthetic 3D worlds require a completely different interaction and navigation paradigm than the classical keyboard-and-mouse approach. This paradigm might even be paired with appropriate visual and environmental clues, such as tactile feedback and 3D sound. Although 3D mice and pointers might seem appropriate for common VR applications, including architectural walk-throughs or terrain visualization, they are not suitable for the envisioned visualization and design tasks. Therefore, the feasibility study is targeting a dual-handed design paradigm in which the user is equipped with a pair of pinch gloves as the main input and interaction device, in addition to the previously

mentioned 3D mice and pointers. The pinch gloves are equipped with basic gesture recognition, which allows the association of natural-hand movements with specific model interactions. Previously complex tasks, such as positioning, rotating, scaling, and modification of objects in 3D space become possible with intuitive hand gestures. This approach allows the preservation of the hands-on experience from the physical world, while overcoming the well known classical 2D constraints introduced with the keyboard.

The project utilizes an immersive workbench, a state-of-the-art VR projection device, that produces table-size, stereoscopic renderings of large-scale 3D data sets. This VR display device allows the projection of 3D computer-generated images onto an approximately 7-foot by 5-foot projection area. Its similarity in appearance to a drafting board brought it early recognition and acceptance by engineers. Even though the user is required to wear a special set of shutter glasses, this environment is semi-immersive. A user interacts with an object that is rendered stereoscopically in 3D space thus, providing the analog of the traditional 3D design space. The scientist can either use data gloves or 3D pointing devices to interact with a virtual model representing either the results of a simulation or its original problem definition.

The user, for example, will be able to grab two points in space and rotate either a selected object or the entire environment around the center point. A truly powerful feature of this implementation is its ability to provide an unprecedented amount of real estate to the user in the form of a 3D desktop. Anyone with exposure to working with multiple open and overlapping windows or virtual 2D desktops on a regular display will appreciate that objects, tools, and other components can now be placed, arranged and viewed in an unlimited 3D domain.

### *Results*

A proof-of-concept framework supporting the mentioned input and output devices was designed and implemented in late summer 1998, and an initial visualization test on a demo data set from LLNL were performed. The user can now visualize an animation sequence of iso-surfaces derived from a 512x512x512 data set, and intuitively inspect it while immersed in a synthetic environment. More advanced support for user interaction was available in the fall.

# Finite Element Modeling in Simulations of Full Flexion of the Human Finger

*Jonathan Pearlman*

*University of California, Berkeley*

## Abstract

The bulk of my time this summer was spent running finite element analyses (FEA) on a model of a human index finger. The finite element model (FEM) previously developed in our Lab includes all of the bones of the finger and many of the soft tissues. Our primary goal this summer was to successfully simulate full flexion of the human finger, defined as a 21-mm excursion of the flexor digitorum superficialis (FDS) and flexor digitorum profundus (FDP) tendons. Our secondary goal was to validate the model by performing simulations of the index finger at postures described in the literature, and comparing joint reaction forces between results from the FEA experiments and values noted in the literature. Simulations addressing both of these goals were run in parallel throughout the five weeks that I worked at Lawrence Livermore National Laboratory.

Index finger flexion simulations were performed under minimal constraints: The metacarpal was fully constrained, with the three more distal bones unconstrained. A displacement of 21mm was prescribed for the most proximal elements of the tendons over a given time span. All simulations were done quasi-statically. Soft tissues (ligaments, annular pulleys, and tendons) were modeled as hyper-elastic material that was developed specifically for simulating this type of biologic tissue. A pre-stretch was prescribed for each of the ligaments, which, during the initial increments of the

simulation, contracted all of the joints so that adjacent articular surfaces were in contact. Following this pre-stretch phase, tendon displacements resulted in flexion of each of the three joints of the finger. The success of a given simulation was judged by the amount of tendon displacement achieved before the simulation failed, and on how equally each of the joints flexed. (If one joint flexed much more than others it was clear that the simulation did not accurately model in-vivo finger flexion.) In the simulations performed this summer, the maximum tendon displacement was 14mm. All simulations failed due to soft tissue inversions (mostly in the tendons.)

To validate the FEM technique, certain postures (joint flexion angles) must be simulated. Using joint angles for a C&T Pinch, from Chao, et al., rigid body rotations were applied to the proximal phalanx, medial phalanx, and distal phalanx. There was limited success in these simulations primarily because of the difficulty of simultaneously displacing the tendons while imposing a rotation on the bones. Invariably, elements in the soft tissues would invert, causing the analysis to fail.

In an effort to improve the convergence of the analysis, the model was continually refined. One of the major refinements added a volar plate on the proximal end of the proximal phalanx. This addition was done to make the ligaments of the metacarpophalangeal (MCP) joint physiologically accurate. The dorsal ligaments of the MCP joint insert into the dorsal plate. Prior to the addition of the volar plate, the dorsal ligaments were attached directly to the proximal phalanx; this drastically reduces the moment arm the ligament has on the bone. The addition of the volar plate, and the subsequent re-meshing of the ligaments of the MCP joint, greatly improved the convergence of the model.

Reference: Chao, E.Y., K. An, W. P. Cooney, and R.L. Linscheid (1989), *Biomechanics of the Hand*, (World Scientific Publishing Co.)



# Testing the ZPL code on ACI Blue-Pacific

*Lyn Reid*

*University of Washington*

## Abstract

My summer project at Lawrence Livermore National Laboratory (LLNL) involved working on neutron transport code with Peter Brown and Anne Greenbaum. Peter is an LLNL researcher in the Center for Applied Scientific Computing (CASC), and Ann, like myself, was also visiting from the University of Washington for the summer. Under her supervision the previous year, I had developed code for one-dimensional steady state neutron transport in a parallel language called ZPL. Developed at the University of Washington, ZPL is a language that ports easily between sequential and parallel machines.

My project began by putting this ZPL code on the IBM ASCI Blue-Pacific machine, and comparing its performance with MPI. This testing uncovered some problems with the ZPL language that will be addressed by the developers. It also served to familiarize me with numerical methods for solving the transport equation and with parallel computing.

With this background, we then began to consider the problem of adaptive mesh refinement in transport codes. We discussed a number of ideas for adaptively refining the mesh or for using an irregular mesh created at the beginning of a run. When working in multidimensions, one must decide how to define the difference scheme at the interface between the underlying grid and a patched zone. Then, one must employ efficient methods for solving the matrix equations that come from such a difference scheme. It is best to retain the ability to march through the grid in the direction of neutron travel, and the difference scheme should allow for this. At the same time, one wishes to take advantage of parallelism as much as possible, and to generate matrix equations for which iterative methods

converge rapidly. With many ideas to consider here, and we did not settle on a single strategy that seemed best. Work in this direction is continuing.

While I was at the Laboratory, I also attended several seminars sponsored by CASC. I had an enjoyable summer, and I feel that I learned a lot in my time spent at LLNL.

# Visualization for SAMRAI Data Sets in AVS/Express

*Michael Wittman*

*University of California, Davis*

## Abstract

Visualization of data created by adaptive mesh refinement (AMR) codes, such as the structured adaptive mesh refinement application infrastructure (SAMRAI) code being developed at Lawrence Livermore National Laboratory, normally is not supported by commercial visualization tools such as AVS/Express. While custom-built visualization tools exist for AMR grids, these do not have the flexibility and rapid development capabilities of AVS/Express. One facet of the SAMRAI Project explores the use of AVS/Express to see whether, with modest development work, it can be extended to meet data visualization needs.

One of the main problems in visualizing AMR data occurs at the boundary interfaces between multiple resolution patches. AVS/Express handles multiple patch data by running the visualization algorithm separately on each patch, and then rendering the individual patches as separate collections of data. This causes problems for techniques that use interpolation methods, such as isosurface and contour creation. Discontinuities are created along patch boundaries when coarse and fine patches are adjacent to each other. The focus of my work was to investigate solutions that reduce these discontinuities.

The first approach explored an unstructured representation and added sutures of additional AVS/Express cells to correctly connect the cells of each of the patches. This added some additional overhead but eliminated the discontinuities, while remaining very close to the original data. After some experimentation, it was

found that this approach would require more time than was available to finish the implementation.

A second, and more traditional, approach was then investigated. Rather than adding cells to create continuity at the patch boundary, the data at the fine level were adjusted to match the coarse level data. This method allowed the standard AVS/Express modules to function on each patch separately, while creating a smooth interface for the isosurface module. Since this can cause some change in the isosurface created, users were consulted to verify that the change would not be significant to their understanding of the data. The second approach was completed and yielded acceptable results. The software developed will be integrated into the SAMRAI visualization toolkit.



# Rayleigh–Taylor Mixing Instability in Miscible Fluids and Layer Formation in Stably Stratified Fluids

*Yuan-Nan Young*

*University of Chicago ASCI Alliance<sup>1</sup>  
Tri-Lab Alliance Technical Sponsor: Stephen Libby*

## Abstract

### Rayleigh–Taylor Mixing Instability in Miscible Fluids

In this work, Yuan-Nan Young studied the Rayleigh–Taylor instability with particular attention to the mixing zone at the interface in miscible fluids. This research was conducted in collaboration with Guy Dimonte, Andrew Cook, and Oleg Schilling of Lawrence Livermore National Laboratory.

An important experimental feature herein addressed computationally was the asymptotic scaling of the mixing zone as  $\Gamma A g t^2$  (where  $\Gamma$  is a constant equal to about .05,  $A$  is the Atwood number, and  $g$  the acceleration). This study focussed both on the accuracy of the scaling law and the dependence of the constant  $\Gamma$  on the form of the initial perturbation. The model system studied was the incompressible Navier–Stokes equation in the Boussinesq approximation (the density contrast for the system causing the Rayleigh–Taylor instability is achieved through linear density dependence on a scalar field). In this case, Atwood numbers are naturally quite small,  $\sim 10^{-3}$ . To model this system, Young modified and adapted a 2D spectral code originally created by Cattaneo, Werne, and Julien. The calculations were done on a 5122 grid. Preliminary calculations for the idealized

case of single modes found growth rates similar those found for immiscible viscous fluids. These studies also demonstrated strong dependence of the early mixing on perturbation wavelength: The higher the perturbation wavelength, the faster the mixing zone grew. Further investigation showed that the mixing zone followed the  $\Gamma A g t^2$  rule independent of initial conditions. However,  $\Gamma$  was discovered to vary from .02 to .06, depending on initial conditions.

### Layer Formation in Stably Stratified Fluids

Layer formation is thought to be an important process in astrophysical phenomena (e.g., semi-convection). A detailed understanding of energy transport and mixing of chemical compositions within layers is therefore highly desirable. Since, from a numerical perspective, it is not yet possible to simulate these phenomena directly, it is appropriate to attempt to capture the physics in a reduced model. One such model is the one-dimensional turbulent model (ODT) designed by A. Kerstein (Sandia National Laboratories/California), which subsequently has been applied to scalar field mixing by Kerstein and S. Wunsch (Chicago, and Sandia National Laboratories/California, respectively).

The problem of layer formation in stably stratified fluids is an example of a test bed in which ODT-based models can be compared both to experimental results and to direct numerical simulation. In this work, Young (in collaboration with Kerstein and Wunsch) took the approach of formulating a toy model to study layer formation in a stirred stratified fluid (a case first studied by Linden and Whitehead). He then applied ODT and direct numerical simulation (DNS) to the model, and compared the results with ODT. Preliminary ODT model and DNS calculations show that, upon adjustment of one free parameter, the ODT model yields agreement with the layer width obtained via DNS (using the same forcing in both cases).

<sup>1</sup> The University of Chicago ASCI Alliance is focussed on advancing the understanding and modeling of astrophysical thermonuclear flashes that are due to the accretion of material onto a compact body and subsequent ignition. These flashes manifest themselves respectively as x-ray bursts, type 1a supernovae, or

novae depending on whether the accretion is onto a neutron star or a white dwarf (with either an off center carbon flash for a supernova 1a or a thin shell hydrogen flash for a nova). These flashes are significant in the production of light elements, in the accurate determination of the cosmic distance scales, and in probing the physics of the compact objects themselves. Since the accumulation of material and its ignition mode depends crucially on turbulent mixing, a significant part of the effort at Chicago is focussed in this area. In particular, the effort of the two LLNL summer students, Jonathan Dursi and Yuan-Nan Young, was aimed at advancing and validating Rayleigh-Taylor interface mixing models, layer formation, and subgrid turbulence models.



# ISCR Annual Report FY 1998

*Lawrence Livermore National Laboratory*  
*P.O. Box 808, L-561*  
*Livermore, CA 94551*

